

THREE KINDS OF PRACTICAL REASON

by

SHANE GRONHOLZ

B.A., Whitworth University, 2005

M.A., University of Colorado Boulder, 2008

A dissertation submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Philosophy

2016

This dissertation entitled:
Three Kinds of Practical Reason
written by Shane Matthew MacDonald Gronholz
has been approved for the Department of Philosophy

Chris Heathwood

Benjamin Hale

Michael Huemer

Alastair Norcross

Graham Oddie

Date 05/04/2016

The final copy of this dissertation has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

ABSTRACT

Gronholz, Shane Matthew MacDonald (Ph.D., Department of Philosophy)

Three Kinds of Practical Reason

Dissertation directed by Professor Chris Heathwood

This dissertation is about the nature and variety of practical reasons. A practical reason is a fact that counts in favor of an agent performing some action, having some emotion, or having some other non-cognitive attitude. I provide a classification of different kinds of reasons and offer an account of how they determine different kinds of *oughts*. Many philosophers recognize that there is something distinctive about moral reasons and the moral point of view, but just what distinguishes the moral from other normative standpoints is not well understood. On my view, what makes a reason a moral reason is that it is essentially other-regarding: it is a reason to treat other individuals in certain ways, for their own sakes. My account of the nature of the moral competes with some of the most prominent normative theories currently on offer, i.e., utilitarianism, Kantian deontology, and virtue ethics. Moral reasons can be further divided into two distinct kinds: general other-regarding reasons and relational reasons. The former are reasons regarding one's treatment of all moral patients. The latter are reasons regarding one's treatment of only certain individuals, namely, those with whom one stands in a special relation, particularly those with whom one has a personal relationship. Non-moral reasons, by contrast, are essentially self-regarding: they are reasons to treat oneself in certain ways, for one's own sake. Both relational reasons and self-regarding reasons are often stronger than general other-regarding reasons, therefore, my view is partial in the sense that an agent is often justified in prioritizing the interests of both herself, and those who are near and dear to her. That there are genuinely normative non-moral reasons may be seen as a threat to the authority many take morality to enjoy. In the final chapter, I explain how and in what sense morality is still authoritative on my view. I do this, in part, by offering a novel account of a moral requirement.

DEDICATION

To Harry MacDonald, for igniting the spark
and to Chris Heathwood, for fanning the flame

ACKNOWLEDGEMENTS

First and foremost, I thank my advisor, Chris Heathwood, for the countless hours he put in to guiding this dissertation. This project is at least as much a product of his efforts as it is of mine; it would have been impossible but for his time, care, and wisdom. There is not a philosopher I admire more, and few people I admire as much, as Chris. I am proud and honored to be his student.

I would also like to thank Ben Hale, Michael Huemer, Alastair Norcross, and Graham Oddie for serving on the dissertation committee. All of their work has played an important role in shaping my views, and their feedback will be vital as I continue to pursue a career in philosophy.

I thank my very dear friend, Barrett Emerick, for his friendship, support, and guidance. He is a better friend than anyone should ever hope or expect to have.

I thank my parents, Marc and Debbie Gronholz, for all the love, encouragement, and so much more that they have given me throughout my life. I am proud to be their son.

Finally, I thank the two most important people in my life: my wife, Stephanie, and my son, Max. Stephanie, you have allowed me to pursue this dream of mine and have supported our family in the process. I have no words to express the amount of gratitude I have for all you have done and the sacrifices you have made. And thank you, Max, for helping me to understand what partiality is all about: my heart grew three times the day you were born.

CONTENTS

INTRODUCTION	1
1. A CLASSIFICATION OF PRACTICAL REASONS.....	10
2. MORAL REASONS AS OTHER-REGARDING REASONS: A DEFENSE.....	35
3. MORAL REASONS AS OTHER-REGARDING REASONS: OBJECTIONS.....	66
4. RELATIONAL REASONS AND PARTIALITY	91
5. SELF-REGARDING REASONS.....	119
6. THE AUTHORITY OF MORALITY	153
BIBLIOGRAPHY	166

INTRODUCTION

1. My Project

This is a project about practical reasons. A practical reason, it is often said, is a fact (or is provided by a fact) that counts in favor of an agent performing some action or having some emotion or non-cognitive attitude (Parfit 2001). When we deliberate about what to do, we are trying to discover what reasons there are for performing the various actions that are open to us, and weighing them against one another. There may be some reason to do option *A*, but there might be more or better reason to do option *B*, and still more reason to do option *C*. When we have most, or best, or decisive reason to perform a given action, that is the action we ought to perform.

But now we might ask, “In what *sense* of ‘ought’ is it true that we ought to perform the action that we have most reason to perform?” When we say that we ought to do what we have most reason to do, are we always saying that we *morally* ought to do what we have most reason to do? We would be if something that moral philosophers often say about moral philosophy (particularly in introductory textbooks) were true, namely, that moral philosophy is the study of what we should do, or how we should live, or what we have most reason to do. These kinds of statements give the impression that morality is a comprehensive guide to action. But there are cases that seem to show that this is not the case. Suppose I am choosing between two stocks and that there is very good evidence that stock *A* will perform better than stock *B*. It would not be odd for my financial advisor to tell me that I ought to purchase stock *A*. But it would be odd to say that I *morally* ought to purchase stock *A*. Purchasing stock *B* might be foolish or imprudent, but not immoral. This suggests that there are different senses of ‘ought’ and likewise different kinds of reasons, since, presumably, what one ought to do is a function of what one has reason to do. If there are different

senses of ‘ought’ and ‘should’, as the above example suggests, a likely explanation for this is that there are different kinds of reasons. For example, it might be that there are moral reasons and prudential reasons, and likewise, moral *oughts* and prudential *oughts*. If something like this is right, moral philosophers have some difficult questions to answer. If there are different kinds of reasons, what are they? What would it even mean for there to be different kinds of reasons? Do different kinds of reasons have different roles or functions? If there are different kinds of reasons, how are they related to one another? What is their structure? How do they determine what we ought to do? If there are different kinds of reasons, and if reasons determine requirements, are there also different kinds of requirements? Is there such a thing as an all-things-considered requirement?

2. Preview

These are the questions I seek to answer in this dissertation. In Chapter 1, I make a case for reasons pluralism – the view that there is more than one kind of reason. I then offer a classification of basic normative practical reasons. On the classification I offer, there are three kinds of practical reasons: Self-regarding reasons, other-regarding reasons, and relational reasons. Self-regarding reasons are reasons an agent has to treat herself in certain ways for her own sake. Both other-regarding reasons and relational reasons are reasons an agent has to treat other individuals in certain ways, for their own sakes. While other-regarding reasons are reasons an agent has to treat everyone in certain ways, in virtue of the kinds of beings they are, relational reasons are reasons an agent has to treat only certain individuals in certain ways, in virtue of a relation in which that agent stands to those individuals. Though there may be a variety of normatively important relations, I will focus on one in particular: that of a personal relationship.

On my view, self-regarding reasons and relational reasons are, all else equal, stronger or weightier than general other-regarding reasons. Thus, the theory of practical rationality I offer here is

partial, in two ways. It allows (and perhaps requires) partial treatment to oneself, and to those near and dear. An upshot of this view is that it can be permissible to provide a benefit to, for example, my own child, at the cost of a greater benefit for a stranger. Likewise, it can be permissible to provide a benefit for myself, at the cost of a greater benefit for a stranger.

Another important distinction I make is that all self-regarding reasons are non-moral reasons. That is to say that reasons an agent has to treat herself in certain ways for her own sake are non-moral reasons. By contrast, all reasons an agent has to treat other individuals in certain ways for their own sakes are moral reasons, thus, both other-regarding reasons and relational reasons are moral reasons. In short, moral reasons are all and only those reasons regarding our treatment of individuals other than ourselves.

The following chapters are devoted to defending this classification. In Chapters 2 and 3, I defend the claim that moral reasons are all and only those reasons an agent has to treat other individuals in certain ways for their own sakes. In Chapter 2, I make arguments in favor of this view. The arguments I make rely on commonsense judgments we make about actions an agent performs that only affect the agent who performs them. Though it is often the case that an agent ought not perform actions that will harm her, it seems quite strange to condemn such actions as immoral. I also offer a variety of theoretical advantages of this view.

In Chapter 3, I defend this claim against a variety of objections. In particular, I focus on a number of potential counterexamples to this view. These include self-regarding racism (racism against one's own race), self-degradation, and what I call "seedy behaviors," which include behaviors such as promiscuous sex, drug use, and gambling. In discussing seedy behaviors, I respond to arguments made by Michael Huemer that suggest that perhaps philosophers and non-philosophers just mean different things by the term 'morality' and its cognates.

Chapter 4 is about partiality toward those near and dear to us. I begin by exploring the very concept of partiality, and explain the ways in which my view is partial. I then provide a novel argument for partiality. I begin by noting that it is appropriate to *feel* a certain way about those near and dear to us, in particular, it is appropriate to have special affection for them, to have attitudes of favoring them. I then argue that it is often appropriate to *act out* of that special affection or favoring attitude, and acting out of those attitudes requires partial treatment.

In Chapter 5, I defend partiality toward oneself. I begin by discussing, and generally endorsing, Roger Crisp's account of agent-relative reasons. Crisp accepts that there are reasons to promote the good, but thinks that in addition, there are agent-relative reasons, grounded in the separateness of persons, that can permit an agent to perform an action that benefits herself, even when she could have produced more impartial good by performing some alternative action. I then defend this account against criticisms made by Brian McElwee. McElwee agrees with Crisp that we are not always required to produce as much impartial good as we can, but he argues that this is not because there agent-relative reasons. Rather, McElwee thinks that all reasons are grounded in impartial value, but we are not always required to do that which we have most reason to do. He also argues that positing agent-relative reasons yields some implausible implications. Contrary to McElwee, I argue that we are indeed required to do that which we have most reason to do, and that positing agent-relative reasons does not have the implausible implications McElwee claims it has.

In the final chapter, I discuss the authority of morality. If there are non-moral reasons that compete with moral reasons, such that an agent is not always required to do that which he has most moral reason to do, then perhaps morality does not have the importance many believe it has. In this final short chapter, I explore one view in particular about the authority of morality. According to moral rationalism, if an agent is morally required to perform an act, then the agent has most reason

to perform that act. I will then show how my view is compatible with this view of morality's authority.

3. Theoretical Commitments

In this project, I will be assuming the truth of some substantive views in metaethics. First, I assume the truth of moral realism, according to which, moral judgments can be objectively true, and when they are, they are true “independently of what any human being, anywhere, in any circumstances whatever, thinks of them” (Shafer-Landau 2003, 2).

Secondly, I will be assuming the truth of reasons externalism. Reasons externalism is a view about what explains the fact that a consideration counts in favor of performing an action or having a desire. It can perhaps best be understood by first contrasting it with its competitor, reasons internalism.¹

Reasons internalists hold that practical reasons must bear some important relationship to features that are, so to speak, internal to the agent to whom the reasons apply, e.g., desires, goals, projects, interests, etc. Bernard Williams calls this set of motivating features a *subjective motivational set*. There are a variety of opinions about just what sort of relationship this relationship is, but many internalists hold that an agent can only have a reason to perform an action if that agent is *motivated* to perform the action, or would be motivated to perform it under certain conditions. But for most, if not all, internalists, that you are currently not motivated to, for example, exercise, does not entail that you have no reason to exercise. This is because you might be so motivated if, for example, you were to be made aware of the benefits of exercises, and perhaps to also vividly imagine what it would be like to have those benefits, and what it would be like not to have them. That you *would* desire to exercise under some relevant (perhaps ideal) conditions can be a reason to exercise, even if

¹ In Chapter 5, I will explore the possibility of voluntaristic normative reasons – facts that we can

you presently lack a desire to exercise. Thus, these views are counterfactual theories of practical reasons. The details about what counterfactuals are relevant for determining a practical reason are, to a large degree, what distinguish internalist theories.

Reasons externalism, on the other hand, holds that practical reasons need not connect in any way to internal features of the agent (e.g., desires, motivations). Instead, reasons have their source outside of the agent's desires or motivations – often value or goodness. According to reasons externalism, I could have a reason to exercise even if I don't want to exercise, I don't care about my health, and I couldn't be made to care about my health. Why? Because exercise is *good* for me. It doesn't matter that I don't care about my health, because I *should* care about my health.

I am of course painting with a broad brush here. My aim is not to provide a detailed account of either view, and much less a defense of reasons externalism.² My aim is only to say that I will be assuming the truth of reasons externalism in this dissertation.

It is worth pausing here to say that internalism and externalism are both theories about *normative* reasons, as opposed to motivating reasons. Of course, for the internalist, there is a very strong connection between motivating and normative reasons. Though it should be noted that according to plausible versions of externalism, desires or motivations are often relevant to practical reasons. If I'm choosing between chocolate and vanilla ice cream, and I have a stronger desire for vanilla, I should, all else equal, choose vanilla, since, because I desire it more, I will likely enjoy it more, and enjoyment is good. At any rate, both internalists and externalists can recognize correct non-normative senses of the term 'reason.'

While very often motivating reasons and normative reasons overlap and there are interesting and important connections between them, the two concepts are distinct. There can be cases where our motivating reasons and our normative reasons come apart. Questions about what motivates a

² For a much more detailed discussion of these views, see Parfit 2011, 43-110.

person are empirical, psychological questions, while questions about normative reasons are philosophical in nature. To see how the two can come apart, consider the following scenario. Suppose I punch someone in the mouth because I do not like the sound of his voice. When asked for the reason I punched him, I might answer, “Because I did not like the sound of his voice.” I have reported what motivated me to act as I did, and what I have said is true. But now suppose someone objects, “But that is no reason at all to punch someone!” That also seems true. But now we have two seemingly contradictory claims: That I both had a reason to punch and I did not have a reason to punch. The only way this can be resolved is if ‘reason’ is being used in two different senses and that is precisely what is going on. This project is about normative reasons.³

The final assumption, or set of assumptions, is about a set of terms I will use often in this project, terms such as ‘reason’, ‘obligation’, ‘requirement’, ‘ought,’ ‘should’, ‘prohibited’, and ‘forbidden’, and their respective cognates. For the purposes of this project, ‘ought’ and ‘should’ are synonymous, as are ‘obligation’ and ‘requirement’, and ‘prohibited’ and ‘forbidden’. Likewise, ‘ought’ and ‘should’ simply mean ‘is required’ or ‘is obligated’. And each of these also mean ‘is forbidden from not’ and ‘is prohibited from not’. Thus, for present purposes all of the following expressions are synonymous.

- i. ‘ x ought to Φ ’
- ii. ‘ x should Φ ’
- iii. ‘ s is obligated to Φ ’
- iv. ‘ x is required to Φ ’

³ Though I am not an internalist about reasons, all of this can be accommodated by an internalist view, since we are sometimes motivated to perform actions that will frustrate our deeper or more long-term goals. For example, I may have a strong overriding goal to not spend the night in jail, and there is a good chance that punching someone in the face will frustrate that goal or desire. Strictly speaking, it would probably be more accurate to say that according to internalism, I have *some* normative reason to punch, but I have most all-things-considered reason not to punch. Though this could be true according to externalism as well, since I might derive some satisfaction from punching.

- v. 's is forbidden from not Φ ing'
- vi. 's is prohibited from not Φ ing'

Furthermore, all of these have to do with reasons. S ought to Φ - S is required to Φ - when S has most reason to Φ . When S has most reason to Φ , S is prohibited from not Φ ing.

While it might be obvious that if S is required to Φ , S ought to Φ , some may balk at the reverse - that if S ought to Φ , S is required Φ . This is because 'required' may sound stronger, or more binding. If you see a frail old woman struggling with her groceries and you have nothing better to do, it may seem natural to say that you ought to help her, but it sounds strained to say you are *required* to help her. But if you find yourself in such a situation, and it is really true that there is no good reason not to help the old woman, you likely have most reason to help her and there really would be something wrong with not doing so. In that sense, you are required to help her. In some cases, as perhaps in this case, you may be required to do something, such that it would be wrong not to do it, but it may not be *very* wrong. In cases where one is inclined to say that S ought to Φ , but is not required to Φ , we can probably remedy this by adding the word *weakly*: S ought to Φ , but is only weakly required to Φ . Both of these claims have to do with what S has most reason to do. I said I will be assuming all of this, but in fact, I will defend the claim that one is required to do that which one has most reason to do in Chapter 5.

Just as there are non-normative senses of the term 'reason,' there are non-normative senses of terms like 'requirement' as well. For example, there are legal requirements that are not necessarily normative requirements. Sometimes we are not *really* required to comply with legal requirements. Suppose I am driving on an empty highway with an unreasonably low speed limit, and that I have very good reason to believe that there are no traffic police monitoring the road. Under these circumstances, it might be that although there exists a legal requirement to obey the speed limit,

there is no real, normative requirement for me to obey the speed limit. It is not important that I follow the speed limit. There is no fact that counts in favor of my obeying the speed limit. I have no reason to obey the speed limit. I may in fact have good reason for disobeying the speed limit.

I could continue to give other examples of requirements that are not normative, for example, those involving etiquette or games. Of course I am not saying that there is no reason to obey the speed limit. Surely there often is such a reason. But the reason is not simply that the law requires it. Rather, the reason will likely depend on the potential harm to oneself and to others.

Having revealed these theoretical commitments, I now move on to offering a classification of practical reasons and defending it.

CHAPTER 1

A Classification of Practical Reasons

The purpose of this chapter is to consider whether there are different kinds of practical reasons, and likewise, different kinds of obligations. Some people are monists about practical reasons. They believe that practical reasons do not come in different kinds. Those who hold this view have a much easier time defending the claim that there are facts about what someone just plain ought to do. If reasons do not vary in kind, there is no apparent difficulty in comparing them, or weighing them against one another. When we posit different kinds of reasons, things get more difficult. If some reasons are of different kinds, it is not obvious that we can compare them and make a determination about what we just plain ought to do in a given situation.

As we have seen in the introduction, I *do* believe that there are different kinds of reasons. In this chapter, I will explain in greater detail what it means for there to be different kinds of reasons, and I will offer a classification of those reasons. This classification will be a *basic* one. There may be instances where it is useful to make a certain distinction among different kinds of reasons, but that distinction itself will be based in a more basic distinction. But here I am only concerned with the most basic kinds of reasons—the distinctions I will be making cannot be reduced to any more basic distinctions.

We have already seen that reasons are provided by facts when such facts “count in favor of having some belief or desire, or acting in some way.” But as Parfit notes, ‘counts in favor of’ just means ‘is a reason for.’ The concept of a reason, then, is fundamental and cannot helpfully be

explained in other terms (2001, 18). Parfit may be right about that, but we can still ask the following question: What do reasons *do*? What is their function?⁴

1. The Function of Reasons

It is not entirely standard practice among philosophers to speak of reasons as having functions but the idea is straightforward. Reasons determine what we ought to do and what we ought not to do. Any reason that is a genuine reason will have, at a minimum, the function of tending to make an act required. We have already noted that a practical reason is a fact (or is provided by a fact) that counts in favor of performing some act. Another way of understanding a practical reason is as a *prima facie* requirement. A *prima facie* requirement is a requirement that can potentially be overridden by some other requirement. If I can provide a benefit to a deserving person, I have a *prima facie* requirement to do so. To see this, consider the claim that, given that I can provide a benefit to a deserving person, and thus have a *prima facie* requirement to do so, it is not the case that I ought to provide the benefit to the individual. In order for such a claim to be true, there would have to be some reason against providing said benefit. Note that such a reason may not be all that significant. It could simply be that providing the benefit would involve some cost to me, or even that I simply don't feel like making the effort (this would be especially plausible if the benefit in question was quite small). But if God informs me that he will see to it that my neighbor experiences an episode of pleasure if I simply tap my nose, we think I should tap my nose, since the fact that my neighbor would experience pleasure is a reason for me to tap my nose. In order to justify my claim

⁴ In asking what the function of a given type of reason is, I do not mean to suggest that reasons are somehow invented. They are not. I simply mean that reasons *do* things, just as water does things but was not invented (e.g., it flows downhill, it boils at 212 °F). Of course, reasons are not physical objects, so they don't do things in quite the same way. The idea of a function of a reason should become clearer below.

that it is not the case that I ought to tap my nose, I would need to point to some other countervailing reason against taping my nose, for example, I believe I may have staph bacteria on my hand that I do not want to spread to my nose, or I believe my neighbor is a bad person and does not deserve an episode of pleasure.

Conversely, a reason not to do something tends to make an act forbidden. When one has a reason not to perform some act, there is a *prima facie* prohibition on performing that act, which is a prohibition that would exist in the absence of any other reasons. For example, if my Φ ing would harm an innocent person, I am prohibited from Φ ing unless there is some other reason *to* Φ . Throwing a rock through my neighbor's window would likely harm her in various ways. Thus, there is a reason not to throw a rock through her window, which is a *prima facie* prohibition on doing so. Given that throwing a rock through her window is likely to harm her in some way, in order to be justified in throwing a rock through her window, I would need some reason for doing so, for example, if her house was on fire and there was someone trapped inside.

Support for this point can be found by considering the standard definition of a reason, i.e., a fact that counts in favor of acting in some way. If there is a fact that counts in favor of your doing something, and no fact that counts against it, then it seems obvious that you ought to do it. If you can provide a benefit to some deserving person by sheer will, and there is no reason not to do it, then you should do it. If someone thinks it is not the case that you should do it, this view can be rational only if there is some *reason* for you not to do it.

In other words, requirements are determined by where the weight of reasons lies. It is both natural, and, I think, appropriate to imagine a scale upon which we weigh our reasons. What we should do is whatever course of action has the weightier reasons. Now, it is true that sometimes the difference in weight will be negligible. In these circumstances, it may sound strange to say that I ought to do what I have the tiniest bit more reason to do, and stranger still to say I am *required* to do

it, since the weight of reasons barely counts in favor of doing one thing rather than the other. But suppose I can very easily provide my neighbor with 1,000 units of value by Φ ing, but can just as easily provide her with 1,001 units of value by Ψ ing. All else being equal, I should Ψ . It might not be very important that I Ψ rather than Φ . It may only be a weak requirement, but it is a requirement nonetheless. Otherwise, I would just arbitrarily be making someone less well off than I could.⁵

It should be noted that not everyone accepts the view that reasons straightforwardly determine obligations. Though it would be difficult to deny that there is *some* relationship between practical reasons and obligations, some deny that reasons *directly* determine obligations in the way I have been describing. Brian McEwlee, for example, claims that “the structure of an obligation is such that we are not necessarily *obliged* to do what which there is most reason to do.” On McEwlee’s view, it can be both true that S has most reason to Φ but S is not required to Φ . In other words, it is sometimes permissible to fail to do what one has most reason to do. Of course, McEwlee does not deny that reasons play an important role in determining obligations. I will discuss McEwlee’s view in detail in Chapter 5. It should be noted however that this view is compatible with the substantive claims I make for the remainder of this chapter.

In general, then, reasons determine obligations, and conversely, prohibitions. I now want to consider the possibility that there are different *kinds* of reasons. There are many ways we might try to classify reasons into different kinds or categories. Some of these differences may not have any normative importance and so the difference in kind will merely be superficial, conventional, or imaginary. For there to be different kinds of reasons, in any normatively significant way, the difference in the kind of reasons would have to make a difference in the normative status of the action that the reasons in question determine. So if I am looking at a pair of reasons and wondering whether they are of a different kind, I should ask whether positing such a difference would make a

⁵ I discuss the possibility of ties in Chapter 5.

difference in the normative status of the action to which those reasons apply. One simple example of a way in which the difference in the kind of reason can make a difference in the normative status of the action to which they apply would be the following: There are moral reasons and prudential reasons and moral reasons always override prudential reasons, such that whenever an agent has any moral reason to Φ , she is all things considered required to Φ , no matter how much prudential reason she has not to Φ . On this (false) classification of reasons, the distinction between moral and prudential reasons would be important, since the difference makes a difference in the normative status of the action under consideration. If someone were blind to the distinction between the kinds of reasons, and only grasped that there was a whole lot of reason not to Φ , and only some reason to Φ , she would determine that she ought not to Φ . But if all of her reasons to Φ were prudential and all of her reasons not to Φ were moral, then, under this classification of reasons, she would be wrong, since she is missing out on the distinction between the moral and the prudential, and on the fact that moral reasons always override prudential reasons.

Contrast this with an example of a way in which the difference in the kind of reason would *not* be important. A simple example of a way in which a difference would *not* be important would be if utilitarianism were the correct view of *practical rationality* (as opposed to just morality). According to utilitarianism, I am required to maximize utility. This theory does not discriminate. Everyone's well-being counts the same. As a view about practical rationality, I only have reason to increase utility (and decrease disutility). One distinction between reasons one might make is between self-regarding reasons and other-regarding reasons. Self-regarding reasons, for simplicity's sake, are reasons an agent has to benefit herself, while other-regarding reasons are reasons an agent has to benefit others. But according to utilitarianism as a view about practical rationality, the difference between these two kinds of reasons would not be important. I have just as much reason to promote my neighbor's welfare as I do my own. Suppose I am considering an act that will provide 10 units of pleasure to my

neighbor. That the reason I have to do this act is an other-regarding reason has no affect on the normative status of the action. All that matters is that *someone* will get the pleasure. If I am obligated to perform this act, I would be just as obligated as I would be if the act would provide *me* with the 10 units of pleasure, and thus my reason to perform the act was a self-regarding reason.

I have said that distinctions between kinds of reasons are normatively important only if positing a difference between kinds of reasons makes a difference in the normative status of the action to which they apply. In the first example, where it did make a difference, the difference in normative status was that between obligatory and not obligatory. But that is not the only way in which normative statuses can differ, for there may be different kinds of obligations. So if action A is x-obligatory, and action B is y-obligatory, though both actions are obligatory, these two actions have different normative statuses. So another way in which the difference between kinds of reasons—which are relevant for determining the moral status of an action—can be normatively important is when the difference itself helps determine what kind of obligation (or prohibition) an agent has. But this will depend on whether there are different kinds of obligations. I will take up this question in the following section.

2. The Case for Reasons Pluralism

I will now make a preliminary case for the view that there is more than one kind of reason. The strategy will be to show that we already have an intuitive grasp of a moral reason. If it can be shown that there are reasons other than moral reasons, then the conclusion that there is more than one kind of reason follows.

The idea of a moral obligation is so familiar it hardly bears mentioning. Virtually any moral realist will agree that we have some moral obligations.⁶ And these obligations are generated by moral reasons. My being morally obligated to save a child drowning in a shallow pond entails, at a minimum, that I have some moral reason to save the child drowning in a shallow pond. To see this, consider the following statement: X is morally obligated to Φ but has no moral reason to Φ . This is an odd statement. How could it be that one could be morally obligated to Φ but have no moral reason to do so?

But now consider a rather different sort of case. One morning, Jason opens his email to find a letter claiming to be from the wife of an important Nigerian official who has now passed away. She stands to receive a large inheritance from her late husband, but to even access the money, there are substantial legal fees she must pay. But since she does not currently have the inheritance, she is quite poor until she gets it, so she cannot pay these fees. She tells Jason that if he sends the relatively small sum of \$5,000 necessary for accessing her money, she will reward him with the handsome sum of \$50,000 as soon as she acquires her inheritance. If the email were true, Jason would not be doing anything illegal or immoral, since neither she nor Jason would be breaking any laws, and she really would be entitled to the money. Moreover, Jason is entitled to do what he pleases with the money; he earned it an appropriate way, he does not owe anyone the money, etc.

Jason is, as we all are by now, well aware of these kinds of email scams. He is not usually one to be taken in by them. Yet, in this case, the email is written so compellingly, that he finds himself actually considering sending the money. He is moved not only by the promise of the financial reward, but he is also truly concerned about this woman's welfare.

⁶ Scalar consequentialists are an exception, but they at least do not deny the existence of moral *reasons*. See Norcross, 2006.

Suppose after some thought, Jason sends the money. Jason has now been successfully scammed. It is obvious that Jason has made a mistake. He should not have sent the money. But when we contrast this with someone who fails to save a child drowning in a shallow pond simply because he does not want to get his clothes muddy, we realize that Jason's mistake is of a quite different sort. Although Jason should not have sent the money—we might say he was rationally forbidden from doing so—he was not morally forbidden from doing so. His failing was not a moral one. He did not do anything morally wrong. While Jason has reason to feel foolish for what he did, a sense of moral guilt would be misplaced. Others might also judge Jason for being foolish, but moral blame or righteous indignation would simply not be appropriate.

If there are moral requirements, and Jason was required not to send the money, but not morally required, then there are different kinds of requirements. If there are different kinds of requirements, and requirements are a function of reasons, then presumably there are different kinds of reasons. If reasons determine requirements, as I have claimed they do, then it is difficult to pinpoint what might differentiate *different kinds* of requirements besides there being different kinds of reasons that ground them.

One might think that Jason *has* in fact done something morally wrong. In sending the money, he is encouraging the scammers to continue scamming. Or, one might argue that although it is true that inasmuch as Jason was moved by pity, his intentions were admirable, he could have done something much better with that money, perhaps giving it to a charity that would use it to save people's lives.

But these explanations for Jason's action being a mistake do not seem to get at the heart of Jason's mistake. Even if it is true that Jason's act was immoral in certain respects, at least one reason why Jason's act was a mistake, and probably the primary reason, is that he threw his money away, and this was bad *for him*. And his harming himself in this way, although certainly an error, does not

seem like a moral error. I could stipulate that there is no chance that his sending the money will encourage the scammers, and that for some reason, there is nothing better he could do with his money, though it is hard to see how this could be the case in the real world.

In thinking about these issues, a perfect case is not readily at hand, since our actions are not performed in a vacuum and they often have ripple effects. There can be many different reasons, and different *kinds* of reasons, that contribute to an act's being required or forbidden, many of which we are often unaware of at the time of our choosing how to act.

We can consider a less vivid example by imagining someone who stays up late sipping whisky and watching reruns of the television show *The Nanny* when he has some important task the following day for which he will need to be well rested. We can stipulate that if he does not do his task well, he will be made considerably worse off, but no one else will be. If he is not well rested, it is unlikely that he will perform his task well. But knowing this, he watches episode after episode as the hours grow late. He keeps telling himself, 'just one more episode,' and convinces himself that he will do just fine tomorrow whether he is rested or not, though he knows in his heart that this really is not the case. It is unlikely that watching reruns of *The Nanny* is of much value at all, either valuable for him or impartially valuable (if such a distinction is genuine). So it is really true that he ought to turn off the television and go to bed. If he stays up and continues watching, he makes a mistake. But again, describing his mistake as a moral one seems inaccurate.

Assuming moral realism (and setting aside scalar consequentialism), we do have moral requirements. But as these cases, and many others like it, show, we have non-moral requirements as well. It is sometimes the case that we ought to do something, but it is not the case that we morally ought to do it. Likewise, we sometimes have a reason to do something, but it is not a moral reason. If that's right, then we must conclude that there is more than one kind of reason. And if the kind of obligation depends on what kind of reasons generate the obligation, then the difference in the kind

of reasons will be normatively important, since an action that is morally required has a different normative status than an act that is non-morally required.

3. A Classification of Practical Reasons

So far, I have claimed that there are different kinds of basic, genuinely normative reasons. I will now say what I think those kinds of reasons are. Borrowing from Douglas Portmore, the first distinction that I shall make is one between a *moral reason* and a *morally relevant reason*. Portmore writes, “I will call any reason that is relevant to determining an act’s moral status a morally relevant reason, and I will call any reason that, morally speaking, counts in favor of, or against, performing some action a moral reason” (2008, 371). This sentence is dense and needs some unpacking. It can seem strange to say that a reason could be relevant to determining an act’s moral status without, morally speaking, counting in favor of, or against, performing that action. But consider an example: Suppose my friend would like a ride to the airport. He could take the bus, but getting a ride from me would be quicker and less expensive. Giving him a ride would be a very nice thing to do. The fact that I would save him some time and money seems to count, morally speaking, in favor of giving him a ride. But suppose I was planning on watching a *Twilight Zone* marathon on television that day. The fact that I would enjoy watching the marathon is indeed a reason to watch it, and it plays a role in determining the moral status of watching the marathon and of giving my friend a ride. Precisely how it determines the moral status depends on the details. If my enjoyment is great enough, and I won’t have a good chance to watch *The Twilight Zone* in the future, it could turn out that refusing to give my friend a ride and watching the marathon is morally permissible; refusing may not be morally wrong. But the fact that I would enjoy watching the marathon does not count, morally speaking, in favor of watching the marathon.

On the other hand, it could be that there is no *Twilight Zone* marathon, and, additionally, that I need to pick up my chainsaw at the shop, and the airport is on the way to the shop. Since I was planning to use the chainsaw to cut down a tree in my yard, I have a reason—on my view it is a non-moral reason—to pick up my chainsaw. Under these circumstances, it probably would be morally wrong to refuse my friend a ride to the airport, since it would make things easier on him, and there would be very little cost to me—I may even prefer his company for that portion of the drive. This again illustrates the way in which the presence or absence of non-moral reasons can help determine the moral status of an action. If the distinction between moral and non-moral reasons is still unclear, I think it will become more so as the rest of the chapter unfolds.

It is also important to note that on my view, moral reasons and merely morally relevant reasons are exhaustive. This is because all genuine reasons are at least morally relevant. A reason's being a moral reason entails that it is morally relevant. In the example above, the fact that I would enjoy watching *The Twilight Zone* is a reason to refuse to give my friend a ride, but not a moral reason. But given that it is a reason, it does play a role in determining an act's moral status. I will explain and defend this point in greater detail in the following chapters. For now, the important point is just that on the view I am defending, only some reasons are moral reasons but all reasons are morally relevant. Since all reasons are morally relevant reasons, in some contexts it will be clearer to refer to reasons that are merely morally relevant but not moral reasons as non-moral reasons, just so long as we keep in mind that, since they are in fact genuine reasons, they are morally relevant.

So just what reasons are moral reasons and what reasons are merely morally relevant? What is the nature of a moral reason? Quite simply, as my examples suggest, I believe that a moral reason is essentially *other-regarding*. On the view I am offering here a moral reason is a reason for an agent to treat other individuals (that is, individuals other than herself) in certain ways, and to have certain attitudes about them, for their own sakes. In particular, they are reasons to treat them well. In the

most basic and straightforward kind of case, other-regarding reasons are reasons we have to benefit other individuals and refrain from harming them. Other-regarding reasons also include the reasons we have to have certain attitudes about others. I am open to the possibility that these are the only basic other-regarding reasons but there might be others. If there are rights, then some of them may be such that we have a reason not to violate them, even if it is possible to violate someone's rights without harming her.

What kinds of individuals do we have reasons to treat in certain ways for their own sakes? Simply, whatever individuals or entities are morally considerable. The view that other-regarding reasons are moral reasons is not committed to any view about which kinds of beings are morally considerable. If, for example, eco-systems or trees are morally considerable, then our reasons regarding our treatment of them are moral reasons. Personally, I am a sentientist about moral considerability, so I will be focusing on sentient beings, and most of my examples will involve human persons. This in no way is meant to suggest that we have no reasons regarding our treatment of animals.

One might also wonder whether we really have reasons to treat all sentient beings well. Maybe some sentient beings do not deserve good treatment, perhaps because they are morally bad. Again, I am neutral on this point. Maybe we do not have reasons to treat undeserving people well. But if we do have reasons to treat them well, those reasons are moral reasons.

Non-moral reasons, on the other hand, are essentially *self-regarding*. They are reasons an agent has to treat herself in certain ways, and have certain attitudes about her self, for her own sake. In the most basic and straightforward kind of case, self-regarding reasons are reasons an agent has to benefit herself, and to refrain from harming herself. They also include the reasons an agent has to have certain attitudes about herself. I am open to the possibility that these are the only basic self-regarding reasons, but there might be others, for example, pursuing various perfectionist values.

The locution ‘for their own sakes’ with respect to other-regarding reasons is important, since sometimes we have self-regarding reasons to treat other individuals in certain ways, for *our own* sakes. On my view, these reasons would not be moral reasons. In addition to the moral reasons I may have to treat my boss well for her own sake, I also have self-regarding reasons to treat her well for my own sake, since she has the authority to fire me or give me a promotion. If I treat my boss well, my chances for a promotion improve. This fact provides me with a reason to treat my boss well, but this reason is not a moral reason. In treating my boss well (in order to get a promotion), I am indirectly treating myself well, for my own sake.

Conversely, we can have reasons to treat ourselves in certain ways, for the sake of someone else. Suppose I am medically healthy but my wife would like me to begin a weight-training regimen, simply because she will find me more attractive if I have toned biceps. Lifting weights can aptly be described as my treating myself in a certain way, but I do it for the sake of my wife. In this case, my wife’s happiness gives me a reason to treat myself in a certain way, for her sake. Thus, I have a moral reason to lift weights (albeit perhaps not an especially weighty one). In treating myself in this way, I am indirectly treating my wife well.⁷

Other-regarding reasons can be further divided into two distinct kinds: *Relational reasons* and *non-relational* reasons. Hereafter, I will call non-relational other-regarding reasons simply ‘other-regarding reasons’ and I will call relational other-regarding reasons simply ‘relational reasons.’ Relational reasons are a species of other-regarding reasons since they are reasons for treating other individuals in certain ways, and are therefore moral reasons. But while other-regarding reasons are reasons we have to treat all beings in certain ways, relational reasons are the reasons we have to treat

⁷ Once again, examples of these sorts are often imperfect, since we often have many reasons, and therefore often different *kinds* of reasons, for performing the same action. I will likely benefit from a weight-training regimen, and thus, on my view, I have non-moral reasons to do so. But my wife’s happiness, for her own sake, is at least one among perhaps many reasons to undergo weight-training.

only *certain* beings in certain ways, namely, those to whom we stand in certain relations, and *because* we stand in those relations. There are basic practical reasons that apply to me and no one else in virtue of my relationship with my wife.⁸

I say ‘basic’ practical reasons because a completely impartial view could accommodate the idea that there are practical reasons that apply to me and no one else in virtue of my relationship with my wife. For example, if my wife and I are having marital problems, I may have a reason to attend marriage counseling with her. Doing this may very well have good consequences from a utilitarian point of view. But it simply would not make sense for someone else to go to marriage counseling with my wife—she is married to *me* after all.

The difference is that on my view, my reasons for treating my wife in certain ways extend beyond those based on consequences. On a consequentialist view, that she is my wife is only relevant insofar as this fact determines the consequences of my action. To use another example, this one quite familiar, suppose there are two women drowning in a lake; one is my wife, the other is a stranger. Suppose I can see into the future and I determine without a doubt that saving the stranger will result in better consequences. According to consequentialism, this would directly entail that I ought to save the stranger. But on my view, it could be the case that I ought to save my wife⁹, even if I know that doing so will ultimately result in a worse state-of-affairs. This is because the fact that she is my wife is normatively significant all by itself.

This concludes the basic classification of reasons. We now have three basic kinds of reasons: other-regarding, relational, and self-regarding. Other-regarding reasons and relational reasons, I have claimed, are moral reasons, while self-regarding reasons are merely morally relevant. The claim that a

⁸ I will discuss how this distinction is normatively important in the section below on the structure of reasons. Here, I simply want to state what I take the basic kinds of reasons to be.

⁹ Though not necessarily. It is also consistent with my view that under certain circumstances, I may be required to save the stranger.

reason is a moral one iff it pertains to how we treat individuals other than ourselves is sweeping and controversial. I will defend this claim in much greater detail in the following two chapters (though the presumptive case for reasons pluralism above involving the email scam provides some clues as to how I will go about doing that). For the remainder of this chapter, I will consider questions about what the claims I have made—that other-regarding and relational reasons are moral reasons and self-regarding reasons are non-moral reasons—amount to. Just what is it for other-regarding and relational reasons to be moral and for self-regarding reasons to be non-moral? What is the structure of these reasons and how do they determine our obligations?

4. The Structure of Reasons

Once again, I begin with (non-relational) other-regarding reasons. Since other-regarding reasons are moral reasons, they generate *prima facie* moral requirements. The reason I have to benefit someone else is a moral reason and it generates a *prima facie* moral requirement to do so. Since it is a *prima facie* requirement, I am not always morally required to benefit someone else, but I am required to do so if there is no reason that pulls in the opposite direction. If I can benefit someone and there is no reason not to, I should do it. Once again, if I can provide some deserving person an episode of pleasure by sheer will, I have a good reason to do so. If there is no reason not to, I should do it.

Relational reasons are also moral reasons and therefore generate *prima facie* moral requirements. My wife and I have basic moral reasons to provide for our children that apply to no one else. To fail to properly provide for our children would be a serious moral failing on our part.

Although other-regarding and relational reasons are both moral reasons, since they are both reasons for treating other individuals in certain ways, there are normatively important differences between them that warrant placing them in separate categories. When only other-regarding reasons are present, we should be impartial. This is not so when relational reasons arrive at the scene. When

a parent has a choice between providing some benefit to her child and providing a greater benefit to a stranger or some set of strangers, it is often permissible, and sometimes obligatory, to give priority to the interests of her children. Suppose Sarah has two children. She could see to it that her and her children's most basic survival needs are met and spend all the rest of her income on famine relief. But I take it that if Sarah did this, most people would think that she was not a very good parent, that she was shirking her duties, and she really ought to do more for her children than just what is necessary to keep them alive. We might admire her generous spirit, but we might think her moral concern is misplaced. Because she is a parent, she really ought to make sure her kids are comfortable, well fed, and are given a good start in life, even if Sarah could save the lives of other children by donating the vast majority of her income to Oxfam. Now, it might be that this predicament, prior to having children, would provide Sarah with a reason not to have children at all. I find questions about whether people should have children at all extremely interesting and important and I do not think the answers to these questions are easy to come by. I will not explore this question any further here. The point is simply that *given* that Sarah currently does have children, she has special obligations in virtue of her relationship to her children.

In this way, relational reasons are weightier than non-relational reasons. The simplest way to see how this might practically play out is to imagine that Sarah has some non-dividable good that she can give to her child or to some other child. Suppose that each child already enjoys the same degree of welfare. Once again, we can imagine a scale. Impartially, Sarah has equal reason to provide the good to each child. But the fact that one of the children is Sarah's own child "tips the scales" in favor of her giving the good to her own child.

Now, I am not saying that as long as Sarah has children, she will always be obligated to spend all of her money on them in favor of donating to Oxfam. It is of course permissible for middle-class parents with extra income to donate some portion of their income to charity, even if it

means, for example, buying their children less or less expensive clothing. At this point, we may find ourselves wondering, “Well, just when should Sarah donate to Oxfam and when should she take her children clothes shopping?” Unfortunately, I see no way of establishing a hard and fast rule to answer these kinds of questions. I can only say that relational reasons are weightier than non-relational reasons. For example, we could say that the reasons a parent has to provide a certain benefit to her own children are always twice the weight as the reasons she has to provide an equal benefit to another child. But I think the reader will agree that such a rule is arbitrary. I’m not saying there is no answer. Perhaps God know the answer. But our moral sensibilities are simply not finely tuned enough to determine precisely how much weightier relational reasons are. I am quite certain that a quarter weighs more than a penny, but I couldn’t begin to guess by how much. Does it weigh twice as much? Three times as much? I have no idea.

Above, I said that relational reasons are weightier than non-relational reasons. How best to understand this claim? It might be like this: Suppose I can benefit a stranger or provide an equal benefit to a loved one. We might say that I have a reason to benefit the stranger, and a reason to benefit my loved one, but the reason to benefit my loved one is a weightier reason. Alternatively, we might say that the fact that the loved one stands in a certain relationship to me provides me with an additional reason to benefit her.

This latter way of understanding the claim may in certain cases lead to the “one thought too many” problem raised by Bernard Williams. Susan Wolf discusses the context in which this problem is raised in the following way: The problem is first raised at the end of “Persons, Character, and Morality” in response to a remark of Charles Fried’s. Fried acknowledges that if a man can save one of two people in equal peril, and one of those in peril is his wife, it would be absurd to insist that he must treat both equally, by, for example, flipping a coin. Fried goes on to suggest a possible justification for this judgment—namely, that ‘where the potential rescuer occupies no office such as

that of captain of a ship, public health official or the like, the occurrence of the accident may itself stand as a sufficient randomizing event to meet the dictates of fairness, so he may prefer his friend, or loved one.’ (72)

Williams responds,

But this construction provides the agent with one thought too many; it might have been hoped by some (for instance, his wife) that his motivating thought, fully spelled out, would be the thought that it was his wife, not that it was his wife and that in situations of this kind, it is permissible to save one’s wife. (18)

Returning to the view that relational reasons are additional reasons, one might worry that it would not be appropriate for an agent to reason thus: I have equal general other-regarding reasons to save the stranger and my wife. But that she is my wife provides me with an additional reason to save her. So I will save her. Instead, one might think it is appropriate that the agent simply reasons: She is my wife, so I will save her.

If one is impressed with this concern, one might then prefer to say that the reason the agent ought to save his wife is simply that she is his wife. But that cannot be right, since he would have some reason to save her even if she were not his wife. Which model then is best for understanding relational reasons?

I propose a way of going between the two horns of this dilemma (if it is indeed a dilemma). It begins by noticing this fact: In order for the agent’s wife to be the sort of being with whom he can have a special relationship in the first place, and thus have relational reasons regarding his treatment of her, she already has to have the features that would make it the case that the agent would have non-relational reasons regarding his treatment of her if she were not his wife. When the agent saves his wife, he does so *because* she is his wife. But her being his wife entails that she has other features that provide him with non-relational reasons to save her, even if she were not his wife. In short,

relational reasons *entail* non-relational reasons. In this way, we might say that the non-relational reasons are subsumed by the relational reasons.

Also, on this model, the difference in the kinds of reasons is normatively important. The agent has a relational reason to save his wife, and this reason is weightier than the non-relational reason to save the stranger, because this reason already has the non-relational reasons built in, but it is weightier because it takes into account the fact that the agent bears a very special relationship to his wife. If there were no such difference in the kinds of reasons at play, then whom he ought to save is underdetermined by the reasons he has. But when we posit relational reasons, which are weightier than non-relational reasons, it becomes true that he ought to save his wife.

Other regarding reasons are moral reasons, and as such they generate moral requirements. But self-regarding reasons, because they are non-moral, do not generate moral requirements (though they can, as it were, extinguish them, as in the example involving the ride to the airport). But they do generate requirements of a kind. We can call them prudential requirements. If I can raise my degree of welfare simply by wishing some benefit for myself into existence, and there is no competing reason not to, I really ought to do that. It would be inadvisable, silly, imprudent, or stupid not to. But, crucially, and as I have already claimed, it generally wouldn't be morally wrong not to.

But while self-regarding reasons are not moral reasons, they are nonetheless morally relevant, so they play a role in determining the moral status of an action. I may have a *prima facie* obligation to donate to Oxfam, but may not all-things-considered be morally required to do so, if, say, it would be too costly for me. Like relational reasons, self-regarding reasons also carry a special weight. Why is this? Why am I sometimes permitted to put my own interests ahead of the interests of strangers, even while recognizing that my interests are not impartially or objectively more important than the interests of strangers? Borrowing from Crisp, I contend that the fact that donating some sum to

Oxfam would generate more value than spending that sum on myself does not conceptually entail that I have more reason to donate to Oxfam. As Crisp writes,

“[V]alues are to be sharply distinguished, conceptually speaking, from *reasons*...A reason...is a property of an action that speaks in favour of or against it. What is a value? A value is a property of something which *makes it good*. The enjoyableness of my playing chess, the fairness of aid programs, and the beauty of the Grand Canyon may perhaps be values...[E]ven if all *values* are to be assessed from the objective point of view, it may be that there are certain agent-relative *reasons* which cannot be captured from or grounded in that point of view. (1996, 60-61)

So, it is permissible (and in some cases, as I will explain in the following chapter, rationally required) for an agent to give priority to her own interests simply because they are *her own* interests. Just as Sarah is permitted, and often required, to give priority to the interests of her children simply because they are her own children, we may also give priority to our own interests simply because they are our own. Again, I will take this question up in more detail in Chapter 5.

Claiming that self-regarding reasons are weightier than other-regarding reasons certainly coheres with common practice, though appealing to common practice only goes so far (maybe it doesn't go anywhere at all). We are concerned with how people ought to behave, not just how they in fact behave. Whether ethics should be partial or totally impartial is of course an age-old question. I will address this question in more detail in Chapter 4. But let me briefly say something in support of partiality here: In addition to the strong intuitive appeal that partial views enjoy, here is another point that counts in favor of a partial view: A strictly impartial view would do serious violence to our judgments about what reactive attitudes toward various kinds of behavior are warranted. Consider Becky. Becky earns \$80,000 per year. She gives one half of her income to charity and lives much more modestly than nearly everyone earning the same amount that she earns. She enjoys some

modest luxuries but her lifestyle is relatively Spartan. I have no doubt that Becky is a very good person (at least in this respect) and she deserves a considerable amount of praise for the choices she has made. And yet, according to an impartial view, and assuming that Becky lives in a city where the cost of living is not well above the average, Becky is not a very good person. She is not discharging her moral requirements. Again, depending on the details, she might even be missing the mark by quite a lot. As long as Becky could get by on \$20,000 or \$30,000 per year, she is subject to some pretty serious moral blame, especially when we consider the lives she could be saving with an additional \$10,000 in donations.

One might argue that the reason we praise Becky is that she is *relatively* much more generous than others, and we want to encourage that sort of behavior, and praising Becky is a way of doing that. This is perfectly consistent with the claim that Becky is still not doing enough. But I am making the stronger claim that Becky really is a good person (at least in this respect) and that she *really is* doing all that morality requires of her. We need not even think of praise and blame as public acts. For present purposes, we can think of praise and blame as private acts—as simply the judgments we form about Becky in our own mind. I may read about Becky in a newspaper and form the judgment, “Becky is a very morally upright person.” I do this not because I think she is better than most but still not morally up to snuff, and I don’t do this because I want to encourage her and others to give to Oxfam. Rather, I think it simply because it seems true to me.

I should point out that my view is consistent with the claim that most of us aren’t doing nearly enough to help the less fortunate. In fact, I think it is plainly true that most of us aren’t doing enough to help the less fortunate.¹⁰

¹⁰ I have asserted that relational and self-regarding reasons are weightier than other-regarding reasons and provided some support for this claim. I will defend this claim and discuss implications of it in greater detail in Chapters 3 and 4.

This distinction between self-regarding and other-regarding reasons I have been discussing is normatively important in the way I have defined, since these reasons generate different kinds of requirements. The difference in the kinds of reasons that are relevant for determining the normative status of an action makes a difference in the normative status of the action. An action that is morally required has a different normative status than an action that is prudentially required.

Not only does the distinction make a difference in the kind of requirement that is generated, it can also make a difference as to whether an act is required at all. On my view, it is sometimes permissible to act in our own interest, even at the cost of making the world objectively worse. So if I must choose between two actions, one which provide me with 500 units of value and another that will provide a stranger with 600 units of value, I may not be required to perform the latter, and this is precisely because of the difference in the kinds of reasons I have to perform these respective actions. So in this way, too, the difference in the kinds of reasons makes a difference in the normative status of the actions to which they apply.

One useful way of thinking about the three kinds of reasons I have posited is to view non-relational other-regarding and self-regarding reasons at opposite ends of a spectrum. Other-regarding reasons are moral reasons and they are impartial. Self-regarding reasons are non-moral reasons and are partial. Relational reasons straddle other-regarding and self-regarding reasons. The property relational reasons share with other-regarding reasons is that they concern how we treat other people, and are therefore moral reasons. The property they share self-regarding reasons is that they are partial.

So between other-regarding reasons, relational reasons, and self-regarding reasons, in what order do we place them in terms of their weight? I have already said that self-regarding reasons are weightier than other-regarding reasons. But the weight of relational reasons depends on the relation in question. Many people would agree that a parent is usually morally required to put the interests of

his children *ahead* of his own, as well as ahead of the interests of strangers. It is less clear that our reasons for benefitting our uncle are weightier than our reasons for benefitting ourselves, but it may be that our reasons for benefitting our uncle are weightier than our reasons for benefitting a stranger. I don't want to give the impression that relational reasons are limited to familial relationships. Far from it. We bear many normatively significant relationships. I may be required to spend an extra hour preparing a lecture for my students, in virtue of my relationship of teacher to student, and this may be so *even if* I could do more good by spending that hour doing something else.

An interesting implication of my view is that in a world that contained only one sentient individual, that individual has no moral reasons to do anything, and none of her actions are morally significant. For this individual, all is morally permissible. I take this entailment to *support* my view, since I expect that many would be inclined to agree that it seems strange to say that such a person could do something morally wrong. If I am the only sentient individual on earth, and I, for example, fail to put the proper care into building my living structure such that I am vulnerable to inclement weather, it certainly seems as though I have made an error, either in practice or in judgment, but it is not immoral to fail to build my living structure properly.

This point also helps us to see the distinction between moral reasons and morally relevant reasons more clearly. If I am the lone survivor after the zombie apocalypse, I have no moral reasons whatsoever, but I still have all manner of non-moral reasons. If I'm living in a bunker and discover a stock of *Twilight Zone* DVDs, assuming I have already taken care of my basic needs, I have very good reason to enjoy myself by watching those DVDs. But the fact that I would enjoy watching *Twilight Zone* episodes is not a moral reason after the zombie apocalypse, neither is it a moral reason in the real world where my friend wants a ride to the airport. But it might (depending on the details) be a

morally relevant reason, since it might make it the case that I am not required to give my friend a ride.

5. Concluding Thoughts

I said at the beginning of this chapter that two desiderata of a classification of reasons is that the reasons posited should be genuinely normative and genuinely distinct. I have indicated along the way how my classification satisfies these desiderata, but it may be helpful to offer a quick review. They are genuinely normative because they really do have normative force and they really do generate obligations. I have relational reasons to provide for my child that really make it the case that I ought to provide for him. I have self-regarding reasons to invest my money wisely that really do make it the case that I ought to put my money in fund A rather than fund B. Again, contrast this with legal requirements. It is sometimes the case that I have no reason to comply with legal requirements.

The distinctions I have made between different kinds of reasons are normatively important because the differences make a difference in the normative status of the actions that they help determine. In the case of the distinction between relational and non-relational other-regarding reasons, this is because relational reasons are weighted more heavily than non-relational reasons. When it comes to the distinction between self-regarding reasons and other-regarding reasons (both relational and non-relational), it is because they generate different kinds of obligations. But we might now wonder what this claim amounts to. I have claimed that there are non-moral obligations that are genuine obligations nonetheless. But if that is right, then a non-moral obligation is still an obligation, with all the normative force that is typically entailed by the term. What then do we gain by saying they are of different kinds?

There are at least two answers to this question. The first is that I think that the question of whether there are different kinds of obligations is intrinsically interesting. The fact is, we have this basic concept of morality, so it is interesting to ask what exactly the content of that concept is. I contend that it has only to do with how we treat others. If that's true, that seems to be an interesting truth.

Second, what kind of obligation is in question will play an important role in determining what kinds of reactive attitudes are warranted when someone complies or fails to comply with her obligations. When someone does something immoral, blame or righteous indignation is warranted. When someone does something merely imprudent, pity rather than moral outrage is warranted. When someone does something that improves her own lot, moral praise is a bizarre reaction. If someone invests wisely in the stock market and enjoys a considerable return on the investment, moral praise seems out of place. Once again, I will discuss these implications in greater detail in Chapters 2 and 3.

I have now offered a classification of the basic kinds of reasons. I have purposefully avoided spending too much time defending this view against some key objections, since I wanted to begin simply by laying out the view. The rest of this project will be discussing the implications of this view and responding to some key objections that are likely to be raised. In the next two chapters, I will defend the claim that moral reasons are exclusively other regarding. In Chapter 4, I will discuss partiality in greater detail, and will defend the way in which my view is partial. In Chapter 5, I will focus on self-regarding reasons, and, for reasons that will become apparent, will defend the claim that we ought to (are required to) do that which we have most reason to do. In the final chapter, I will explore whether my view is a version of, or is at least compatible with moral rationalism, the view that one has most reason to do that which one is morally required to do.

CHAPTER 2

Moral Reasons as Other-regarding Reasons: A Defense

In the previous chapter, I claimed that moral reasons are reasons we have to treat other individuals in certain ways for their own sakes and that an agent never has moral reasons to treat herself in any way, for her own sake. The aim of this chapter is to defend those claims.

I will begin by noting that I do not believe it is possible to prove deductively that moral reasons only involve our treatment of others. While I nevertheless think it is reasonable to accept this claim and will attempt to show why this is so, I will first discuss some attempts at proving that agents do not have moral duties to themselves, and explain why I think these attempts fail.

1. Of what sort of proof is the claim that morality only concerns our treatment of others susceptible?

I know of no way of proving that moral reasons only involve our treatment of others. If it is true, it is a brute normative fact, and these cannot be proven directly with a syllogism. Marcus G. Singer, and more recently Joshua Glasgow, have attempted to show that the concept of a (moral) duty to oneself is incoherent. They argue that if A has a duty to B, then B can release A from that duty at any time. But if A has a duty to A, then A can release A from that duty at any given time. But this seems to violate our intuitive notion of a duty. A duty from which an agent can release herself whenever she chooses isn't much of a duty at all. In Singer's words, "It is essential to the nature of an obligation that no one can release himself from an obligation by not wishing to perform it or by deciding not to perform it, or, indeed, in any other way whatsoever" (1963, 202). Glasgow puts it this way:

There seems to be a presumptive case for the claim that (1) any agent can, at her discretion, release herself from any putative self-regarding duty. But, (2) if it is purely a matter of discretion whether an agent performs some action, then that action is not actually obligatory, so (3) there are no duties to self.

The problem with this line of reasoning is that many of our duties to others are not such that they can release us from them. If A has a duty to B, it is not always the case that B can release A from that duty at any time. If I see a man drowning in a lake, it is likely that I have a duty to save him. Suppose upon diving in and reaching the man, he informs me that he does not deserve to live, since he has had too many lustful thoughts about women. He therefore releases me from my duty to save him. It is plausible that I still have a duty to save him. So many duties to others are not, in fact, *releasable*. If it is true in general that not all duties are releasable, then it could very well be the case that duties to oneself are not releasable, and if that is right, then a non-releasable duty to oneself will jibe with our intuitive notion of a duty, that is, duty in the robust sense of the word.

Of course I do not deny that some duties are releasable. Many of them probably are. If I owe my friend \$20 and he tells me, “Don’t worry about it,” he has probably released me from my duty (though I may be obliged to push back a little). Many promises are also such that the promisee can release the promiser from her duties as promiser (though perhaps not always). I am simply denying that duties in general are releasable.

Singer presents a further argument for this point. He asserts, “in general, if A has a duty to B, then B has a right against or with respect to A.” But if I can have a duty to myself, then that entails that I have a right against or with respect to myself. “And this,” says Singer, “is surely nonsense. What could it mean to have a right or a claim against oneself? (Could one sue oneself in a court of law for return of the money one owes oneself)” (202).

I agree that it is bizarre to claim that one could have a right against oneself, but why believe that “if A has a duty to B, then B has a right against or with respect to A”? Is this supposed to be obvious? Or analytic? There are counter-examples to the claim that duties entail corresponding rights. Carl Cohen discusses one such counter-example, when he asserts that he may have a duty to help pay for his son’s college, but his son may not have a right to such assistance (1997, 94).

Singer, aware of such counter-examples, writes,

My answer to this objection is that the relation I am presupposing as holding between rights and duties...is only asserted as holding for those duties or obligations that one can be said to have to someone... It is important here not to confuse the person to whom one is under an obligation with the person regarding whom one is under the obligation - that is, the person who stands to benefit by its performance - for they are not always one and the same. (204)

So perhaps we should say that a father has a general duty to promote the welfare of his children, or help them satisfy their desires, or to promote their education, or whatever. But saying that Cohen has a duty *to* his son misdescribes what is going on.

My response to this is that it seems perfectly appropriate to say that Cohen has a duty *to* his son. In fact, many of the duties a parent has to his child are such that the child may have no corresponding right against the parent, and yet we very often talk about parents’ duties to their (specific, actual) children. And not only do we speak that way, it really seems true. If I can afford it, I may have a duty to my son to send him to a good private school (perhaps if the local public schools are not very good), but if I fail to do that, I have not violated any right of his.

Alternatively, one could simply turn Singer’s *modus tollens*:

P1: If A has a duty to B, then B has a right against or with respect to A.

C1: If I have a duty to myself, then I have a right against or with respect to myself.

P2: But I have no rights against or with respect to myself.

C2: Therefore, I do not have a duty to myself.

into the following *modus ponens*:

P1: If A has a duty to B, then B has a right against or with respect to A.

C1: If I have a duty to myself, then I have a right against or with respect to myself.

P2: I do have duties to myself.

C2: Therefore, I do have a right against or with respect to myself.

If one is already attracted to the idea that one can have duties to oneself, I see no reason not to make this move. The claim that one has duties to oneself seems at least as plausible as the claim that one cannot have rights against oneself. This is especially the case if the claim that “if A has a duty to B, then B has a right against or with respect to A” is intended to be analytically true. One might ask, ‘What would it mean for me to have a right against myself besides simply that I have a duty to myself?’ In other words, if one is already relatively certain that one can have duties to oneself, she may simply take Singer’s argument as a reason to accept that one can have a right against oneself. As I mentioned, I think it sounds strange to say that one can have a right against oneself, but I also think it is strange to claim that one has a duty to oneself. If we do have duties to ourselves, and duties entail rights, then it simply follows that one can have a right against oneself.

So I have doubts that the claim that morality does not involve an agent’s treatment of herself (and related claims) can be proven. But while I cannot prove it, I will here present several considerations that support, motivate, or make plausible this claim. If the view can be adequately motivated, and there are no compelling counter-examples, the view can be rationally held.

2. Motivating the View

In order to motivate the claim that morality only has to do with our treatment of others, we may begin by considering a number of cases. We can begin considering many cases where it seems true that an agent has a reason to treat another individual in a certain way. Call all cases of this type *other-regarding* cases. Then we can consider cases where it seems true that an agent has a reason to treat herself in a certain way. Call all cases of this type *self-regarding* cases. If, upon reflection, it consistently seems that the reasons involved in other-regarding cases are moral reasons, and the reasons involved in self-regarding cases are non-moral, it will be reasonable to form the hypothesis, however provisional, that all other-regarding reasons are moral reasons, and all self-regarding reasons are non-moral. If we come across cases that conflict with this hypothesis, we will have to reconsider. But I do not think there are any compelling cases that show that one has a moral reason to treat oneself in a certain way, or that one's reasons regarding one's treatment of others are non-moral.

Before considering specific cases, I want to point out a general phenomenon. When philosophers are in need of an example of an act that is paradigmatically wrong, perhaps as a counter-example to some normative theory, it inevitably involves one person's mistreatment of another person. Though this of course does not *prove* that the domain of morality is limited to our treatment of others, it is telling. We might ask why the example never or seldom involves an agent harming herself. The answer, I believe, is that it is just much less clear that these types of actions are morally wrong.

I have already briefly made reference to a kind of test for whether a reason, or an obligation, is a moral one. In general, if one fails to respond appropriately to one's moral reasons, the appropriate reactive attitude on her part is guilt, on the part of others, blame or righteous indignation. As Mill puts it, "We do not call anything wrong, unless we mean to imply that a person

ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow creatures, if not by opinion, by the reproaches of his own conscience” (2001, 47). In John Skorupski’s discussion of this passage, he proposes that we “drop the reference to law, and concentrate on the reproaches of conscience and the opinion of our fellow creatures.” He explains, “We lose nothing by doing so, and we leave space for the possibility that something not morally wrong should nevertheless be punished by law” (1999, 141). I concur.¹¹

But blame on the part of others and the reproaches of the agent’s own conscience are not appropriate when the agent fails to respond appropriately to her self-regarding reasons. As Richard Joyce argues,

Our basic emotional response to someone’s self-harm is *pity*. The emotion of *retributive anger* makes little sense within the frame work of prudential normativity...Harming oneself per se doesn’t (and shouldn’t) provoke the emotion of *guilt*; it provokes the phenomenologically very different form of self-castigation of thinking “I’m stupid.” (2008, 63)

It will be useful to keep this test in mind while considering cases.

When I say that blame is the appropriate response when someone else fails to appropriately respond to her other-regarding reasons, the blame does not have to be what Angela M. Smith calls “active blame” (2007). Smith makes a useful distinction between a person’s being responsible and our holding them responsible. She believes that it is possible to believe that a person is morally responsible without holding her morally responsible. She writes,

“To say that a person is morally responsible for something is to say that it can be attributed to her in the way that is required in order for it to be a basis of moral

¹¹ Strictly speaking, based on the quote, Mill is not committed to the view that we should only punish people who have done something morally wrong. My main point is simply that I am going to focus on reactive attitudes, of the agent herself and of others, rather than legal punishment.

appraisal...To say that a person is morally responsible for something, then, is merely to say that she is connected to it in such a way that it can, in principle, serve as a basis for moral appraisal of that person...When we say that we ‘hold’ someone morally responsible for something she has done, for example, what we usually mean is that we *blame* her for it.” (467-468)

When we “hold” someone morally responsible, “we often feel emotions such as anger, resentment, or indignation toward her as a result” (477).

It is important that on Smith’s view (with which I agree), neither judging that a person is morally responsible nor “holding” her responsible requires any expression of the reactive attitude in question. They don’t require that we actively *do* anything, though often we may publicly censure someone whom we hold responsible. It is perhaps obvious that believing or judging that one is morally responsible can go unexpressed, but less obvious that active blame or holding responsible can go unexpressed. But on Smith’s account,

“one can ‘actively blame’ a person simply by feeling resentment, indignation, or anger toward her, without ever expressing these emotions in any way. What unites these unexpressed attitudes and more explicit sanctioning activities, in my view, is that they both go beyond the mere judgment of culpability insofar as they involve the adoption of particular negative attitudes, the expression of criticism, or the performance of sanctioning behavior toward the wrongdoer. It is these further responses that we forswear when we forgive someone, and it is these further responses that we disclaim when we say such things as, ‘I know what he did was wrong, but I do not blame him.’” (477)

Smith points to at least three reasons why we might not hold someone morally responsible, or actively blame her, for something, even though we recognize that she is morally responsible for it.

The first involves standing. She writes, “In some cases, it is simply not our place to reproach another for a faulty action or attitude, even if we think such reproach is deserved” (478). Our relationship to the agent in question can be relevant to determining our standing, as can our own moral character, since it is often not appropriate to be righteously indignant toward someone for a fault that we ourselves possess, even if we think this fault indeed reflects poorly on her moral character.

The second mitigating factor she mentions is the significance of the fault. Relatively minor moral infractions may warrant our modest moral disapproval, but they do not reach a threshold where active blame or reproach becomes appropriate.

Finally, the agent’s response to her infraction can be relevant to whether it makes sense to hold her responsible. If the agent is already remorseful and/or taking steps to correct the wrong, it will often not be appropriate to actively blame her.¹²

The weaker kind of blame, call it “passive blame” is all that is required for my test. When an agent does something that harms only herself, we do not think that this, to use Smith’s language, reflects poorly on her morally, or is personally morally discrediting. When I am asking whether it is appropriate to blame someone, I do not necessarily mean it in the strong sense of actively blaming her. It can be in the weaker sense, that is, in the sense that we merely believe that she is morally responsible. It need not be a strong emotional reaction.

I take it that the claim that self-regarding reasons are non-moral is more controversial than the claim that other-regarding reasons are moral reasons. If I have decisive reason to treat my neighbor in a certain way, for his own sake, and I fail to do that, it seems like that will always be morally wrong. If I owe someone money and have reasons to pay him back for his own sake, and

¹² There are probably additional reasons why one ought not actively blame someone even if that person has done something morally wrong. For example, it can take a toll on the blamer. In many cases, we are psychologically better off if we let go of our resentment, even if the wrongdoer feels no remorse for what she has done. Thus, it may be more precise to say that when an agent has done something morally wrong, it is *prima facie* appropriate for others to blame her.

these reasons make it the case that I ought to pay him back, then it would be morally wrong not to do so. So I won't consider instances of other-regarding cases in detail, except to contrast them with instances of self-regarding cases. There are things we do for ourselves that would be utterly bizarre to think were morally good or praiseworthy, but when those same actions are done for other people, it becomes entirely appropriate to praise the action or to think well of the person's character who performed it. Suppose I go down to the expensive department store to buy myself a new, fancy winter coat. I've already got one that will do – I just want something new. No one thinks there is anything morally praiseworthy about that. But now suppose on my way out, I encounter someone who appears to be down on his luck. He might not be homeless – it could just be that I see that he is less well off than I am. I notice he already has a warm coat on, but it is quite old and shabby. Suppose I give him my new, expensive coat. This would be a terribly kind and generous thing to do. It would be morally praiseworthy. But why such a stark difference in our moral evaluations between keeping the coat for myself and giving it to the stranger? He already has a coat, so neither of us needs the coat for survival. Assume we would each get roughly the same benefit from the coat. Either we are just very confused in the way we evaluate these kinds of behaviors, or there is something fundamentally different about these two actions, that justifies our giving them different moral evaluations. I believe that the best explanation for the difference in moral evaluation between keeping the coat for myself and giving it to a stranger is that, although I have reasons for keeping the coat for myself, these reasons are non-moral, while the reasons to give the jacket to the stranger are moral reasons.¹³

¹³ One worry here might be that I have cheated in my example by making the stranger less well off than I am. Perhaps what is generating the intuition that there is something morally better about giving the stranger the coat is the thought that he would benefit more from it than I would due to diminishing marginal utility. I have three responses to this worry. First, I stipulated in my example that we would both benefit equally from it, and it isn't implausible that, given various details, we would benefit equally. Second, the claim is not merely that it would be morally *better* to give the coat

I have already discussed some instances of self-regarding cases. In the previous chapter, I discussed the case of Jason, who gets taken in by a Nigerian email scam. There, I was attempting only to show that there is more than one kind of reason, since Jason certainly has a reason to refrain from wiring the money, but this reason is non-moral. But now I intend this case to be one among many cases we can consider with the aim of forming a hypothesis that all self-regarding reasons, and obligations, are non-moral. Applying the above test, it does not seem appropriate for Jason to feel a sense of moral guilt, and this is precisely because he has only harmed himself. For some, in some cases, the test may be superfluous. It may already be obvious that the reason in question is not a moral one. But the test may be helpful in harder cases where our intuitions are less clear.

There is also the case where I stayed up drinking and watching “The Nanny” when I knew that this would prevent me from completing an important task the following day, and where this would only harm myself. Once again, it is strange on its face to say that my reasons for turning off the television and going to bed are moral, and stranger still to say that it would be appropriate for a third party to blame me morally for not doing so.

And there are of course many, many more cases, both hypothetical and actual, that we might continue to consider. Describing case after case would be a tedious exercise so I won’t continue it any further. All cases will have the same structure: They will be cases where an agent has a reasons to perform an action, or the agent ought to perform the action, and no one besides the agent herself

to the stranger. It is the stronger claim that there is absolutely nothing praiseworthy about keeping the coat for myself, and I have no moral reasons for doing so. So even if it is true that the stranger would benefit *more* than I would, that is not enough to explain the very stark difference in our judgments of the respective actions. Of course, utilitarians will claim that I *do* have moral reasons for keeping that coat for myself, since I will benefit from having it. But as I have claimed elsewhere, this implication of utilitarianism is implausible, and should be taken as a reason to *reject* utilitarianism. A view that does not have this implication is all the better for it. Finally, the stranger need not even be worse off than I am for my action to be morally praiseworthy. That detail was only added to explain a psychological motivation for giving the stranger the coat, since it would be somewhat puzzling for me to give the coat to someone who was equally or more well off than I am. But most of us would still consider it a very nice (read: morally praiseworthy) thing to do.

will be affected by the action. I contend that in all such cases, it will seem odd or strained or simply false that the agent's reasons to perform the action are moral reasons, or that the agent morally ought to perform the action.

One complication in considering such cases is that very often the relevant actions affect others, and therefore have moral implications. So lending support to this idea is not as simple as providing a list of examples of various actions one can take toward oneself, and then pointing out that in all cases, it seems as though all the reasons the agent has to perform the actions are non-moral. For example, it won't work to say, "My reasons for getting a good night sleep, eating well, drinking coffee, going to work, getting an education, starting a small business, are all self-regarding reasons, and don't these seem to be non-moral in nature?" When I am well rested, fed, and I've had my coffee, I am more likely to be productive and to have a cheery disposition toward others. When I am educated, I am much better suited to make the world a better place. If starting a small business contributes to the overall flourishing of the economy, I may have moral reasons, in addition to the non-moral reasons, to start a small business. Even in the case of the email scam, which I intended to be an example where Jason's reasons for refraining from wiring the money were straightforwardly self-regarding, he may, as a consequence of wiring the money, be unable to pay his friend some money he owes him. If that were the case, Jason would in fact have a moral reason to refrain from wiring the money.

But this complication can be overcome in one of two ways. The first is simply to stipulate, as I have done in describing these cases, that only the agent in question will be affected by the action. If we stipulate that, we avoid the complication of how others will be affected by the action, which can muddle our intuitions and tempt us into thinking that perhaps the agent's reasons are moral, or that the agent morally ought to perform, or refrain from performing, the action in question.

The other way the complication can be overcome is by separating out the various reasons the agent has. If going to bed at a reasonable hour will allow me greater enjoyment of the following day, I have non-moral reasons to go to bed at a reasonable hour. If it will also make me more productive when I spend the afternoon volunteering for Habitat for Humanity, I have moral reasons to go to bed at a reasonable hour.

In this section, I have tried to motivate the view that only reasons for treating other individuals in certain ways are moral reasons. Specifically, I have done this by considering cases where an agent fails to appropriately respond to her self-regarding reasons, and have pointed out that in such cases, blame or righteous indignation seem misplaced. Since I have now made several appeals to reactive attitudes like praise and blame, here is a good place to say something about them in more detail and discuss how they fit into my overall view.

3. Reactive Attitudes

I have claimed that actions of a certain kind make appropriate certain reactive attitudes. For example, when someone does something morally wrong, it is generally appropriate that she feel guilty, and that others blame her (though, as I have discussed, there are exceptions). I have used this point to argue that when an agent fails to properly respond to her self-regarding reasons, and her action only harms (or involves) herself, it is not appropriate for her to feel guilty or for others to blame her. This means that her action was not in fact morally wrong (though it could still be that she ought not to have done it), and that her reasons to refrain from performing the action were not moral reasons.

Let me now address reactive attitudes in more detail. First, what exactly is a reactive attitude? Reactive attitudes are attitudes an agent takes up in response (or in reaction) to some behavior or attitude – either that of another, or her own. They are mental states, which include beliefs, but they

also include emotions, which can “go beyond” mere belief. As Hurley and Macnamara write, “it is one thing to *believe that* someone has slighted you and quite another to *resent* her for the insult” (2011, 373). Often these go together, but they don’t have to. As I am construing them, a mere belief that someone has done something morally wrong counts as a reactive attitude. As does a feeling of anger one might have toward someone else, while recognizing that the anger is misplaced since the person did not actually do anything to warrant anger. Here, the paradigmatic reactive attitudes are praise and blame, which are often the combination of a belief and a retributive emotion (anger, resentment, gratitude).

Sometimes our attitudes can’t be so tidily categorized as either a belief or an emotion. For example, Tamar Szabó Gendler has introduced the concept of an *alief*. “A paradigmatic *alief*,” writes Gendler,

is a mental state with associatively-linked content that is representational, affective, behavioral, and that is activated—consciously or non-consciously—by features of the subject’s internal or ambient environment. Aliefs may be either occurrent or dispositional. (2008, 624)

So, for example, a person may believe that black people are no more likely to mug her than people of any other race, but may still clutch her purse when she sees a black man coming toward her. I won’t dwell on aliefs here. My point in mentioning them is just to acknowledge that there is an enormous spectrum of very fine-grained mental states that are not easily captured by our language. While I will be focusing primarily on beliefs and emotions in my discussion of reactive attitudes, there may be others.

To claim that certain reactive attitudes are appropriate in light of certain actions raises some questions, which may be potentially troublesome for my view. What does it mean to say that a certain attitude is appropriate, or warranted, or makes sense? Presumably, it means one has a reason

to have such an attitude. But then, what kind of reason would that be? I have been claiming that practical reasons divide into three kinds: Other-regarding, relational, and self-regarding. I have further claimed that the former two kinds of reasons are moral reasons, while the latter are non-moral reasons. So my view is in trouble if the reasons we intuitively think we have for having certain reactive attitudes don't match up with the account I have offered. But, as I will attempt to show, I think my classification matches up well with what we intuitively think are the reasons for holding various reactive attitudes.

A clarification: So far, I have been speaking of practical reasons as reasons for performing certain actions or to treat individuals in certain ways. This, as will now become clear, is incomplete. For on my view, one can also have reasons to take up certain attitudes about individuals (others or themselves) and their actions. I speak mainly in terms of action and treatment for the sake of convenience. It could be that taking up a certain attitude about someone or her actions is a way of *treating* her. I'm not certain what my linguistic intuitions are in regard to this matter. If taking up an attitude can be a way of treating someone, fine. If the reader thinks this is odd, the reader should understand that when I talk of practical reasons as being reasons to perform an action or to treat individuals in a certain way, they can also, on my view, be reasons to have certain attitudes, even if that is not always explicitly mentioned.

To see whether my classification matches up with what we think are the reasons for having various reactive attitudes are, I will now consider four action/reactive attitude pairs:

1. The reactive attitude an agent has reason to have when he fails to properly respond to his other-regarding reasons
2. The reactive attitude an agent has reason to have when he fails to properly respond to his self-regarding reasons

3. The reactive attitude others have reason to have when an agent fails to properly respond to his other-regarding reasons.

4. The reactive attitude others have reason to have when an agent fails to properly respond to his self-regarding reasons

1. I will begin by considering the reactive attitude an agent ought to have about his failure to properly respond to his other-regarding reasons. When an agent fails to properly respond to his other-regarding reasons, he has, on my view, done something morally wrong, and therefore, it is appropriate that he experience guilt. Presumably, if it is appropriate for him to feel guilty, then he has a reason to feel guilty. And intuitively, this is a moral reason. After all, it is the appropriate response to an immoral action and is therefore in the moral realm. More importantly, we often think there is something morally wrong with someone who does not feel guilty for having committed an immoral act. The potential danger to my view would be if this attitude is not other-regarding, and therefore, not a moral reason. And one might suppose that it is not an other-regarding reason. After all, guilt is a private feeling, or judgment, or belief, in response to what a person has done himself. But I am claiming that only other-regarding reasons are moral reasons. So how could a reason to feel guilty constitute an other-regarding reason?

Well, on my view, one acts immorally only when one fails to appropriately respond to one's other-regarding reasons. Suppose Timothy did something morally wrong. This can be true, on my view, only if Timothy failed to properly respond to his other-regarding reasons. Suppose he has decisive reason not to steal Tabitha's lunch, but he does it anyway. He has now done something morally wrong, and he therefore has reason to feel guilty. What is that reason? It is that he has failed to treat Tabitha the way she ought to be treated. This reason is straightforwardly an other-regarding reason. If Timothy felt no guilt for stealing Tabitha's lunch, we would think he was failing to properly regard her. Tabitha is a moral patient. She also happens to be a person, and there are ways

to properly regard persons and ways to fail to properly regard them. The notion of respect comes to mind. In general, people ought to be respected, not just in one's actions, but in one's thoughts and attitudes. It is not my aim here to give a full account of how we ought to regard others and why. I will however note here that I do think there are non-consequentialist reasons to regard others in certain ways, perhaps to respect them. It seems to me that racism is wrong even if no one is harmed by someone's racism. For present purposes, it is enough to say that there is a proper way to regard other individuals. I just said that I think there are non-consequentialist reasons to regard others in certain ways, but this does not entail that the view I am defending here, i.e., that an agent's reason to feel guilt when he wrongs others is a moral reason, is inconsistent with consequentialism. That is, we could give a consequentialist account of the wrongness of failing to regard Tabitha in the way she ought to be regarded. Very often, the attitudes we have about people affect how we treat them. If Timothy cultivates respect for other people in general and Tabitha in particular, he will be more likely to treat them well. If he fails to do this, it may make him more likely to fail to treat them well. Therefore, Timothy morally ought to try to cultivate respect (or something like it) for other people in general, and for Tabitha in particular.

The important point, that I suspect many will agree with, is that there is *some* way to properly regard persons, and Timothy has failed to properly regard Tabitha when he feels no guilt for stealing her lunch. His failure to feel guilty for what he has done demonstrates a kind of contempt for Tabitha, and suggests that he does not believe her interests matter. When Timothy feels no guilt for stealing Tabitha's lunch, he is morally in the wrong because he fails to properly respond to his other-regarding reason to regard Tabitha in a certain way. A reason to regard another individual in a certain way is straightforwardly both other-regarding and moral. This result confirms my view.

But our reasons for feeling guilty for wronging other individuals need not be limited to persons. Whatever kinds of beings we take to be morally considerable, it seems as though there will

be a corresponding way to properly regard them. Again, that might be respect, but it need not be. Perhaps only persons deserve respect, but persons are not the only morally considerable beings. If bears are morally considerable, and I wrong a bear, but I fail to feel guilty, the wrongness of not feeling guilty is in my failing to properly regard the bear, whatever the proper way to regard a bear might be.

This is not an *ad hoc* response to this problem. This seems to explain precisely what we think is wrong with someone who feels no guilt when he has wronged someone else. If we point out to Timothy that he did something wrong by stealing Tabitha's lunch, he might respond, "You're right! I feel awful. I knew I shouldn't have done it. It was Tabitha's lunch after all. But I forgot my lunch, and I was so hungry, and I had a moment of weakness. I will go apologize to her." This would be an appropriate response. Although it was wrong to steal Tabitha's lunch, the wrongness ends there. If, on the other hand, Timothy were to shrug his shoulders and say, "Yeah, so what?" we would think that in addition to the wrongness of stealing the lunch, Timothy *remains* in the wrong. But why? The answer seems to be precisely that Timothy fails to show proper regard (respect) for Tabitha. He cannot be bothered to feel any kind of remorse. To Timothy, Tabitha is not even the kind of being to whom any regard is due.

2. When an agent fails to properly respond to his self-regarding reasons, what are the appropriate reactive attitudes on his part, and what reasons are there for those reactive attitudes? In general, it is appropriate for him to (non-morally) regret what he did, that is, to wish he had not done it. Or to think "I'm a fool!" He might also have reasons to feel embarrassed about what he did. Are the reasons for these attitudes moral reasons? Intuitively, no, and that is precisely what is entailed by my view, since, on my view, moral reasons are other-regarding, and non-moral reasons are self-regarding. When an agent fails to properly respond to his self-regarding reasons, he only makes himself worse off, and his attitude is about himself and his imprudent action, and no one else.

In failing to appropriately respond to his non-moral reasons, he has not done anything morally wrong, so on my view, it would be odd to think he has a moral reason to take up a reactive attitude in response to an action that itself is not morally wrong. And this is precisely what intuition tells us. As I've pointed out, a sense of moral guilt would be misplaced here. So what reason does he have to take up these attitudes? I said above that I believe there are non-consequentialist reasons to regard others in certain ways. Likewise, there may be non-consequentialist reasons for an agent to regard himself in a certain way. But whereas our putative non-consequentialist reasons to regard others in certain ways are moral, the reasons an agent has to regard himself in some way are non-moral. To see this, consider a person who falsely believes that she is no good, or that she deserves to suffer, or hates herself without good reason to. She is failing to properly respond to reasons, but her failure does not seem worthy of blame on the part of others. If we should feel anything about her, it should probably be pity or something similar. Though her reasons to regard herself in a certain way are non-moral, they are reasons nonetheless. We should care about ourselves, and when an agent does something that harms herself, she should care about that, and she should care about it in a certain way. That is, she should care that it was *she* who harmed herself. (Similarly, I have reason to feel concerned when someone else is made worse off in some way. But I have a special kind of reason when it was *I* who made her worse off.)

But as before, we can give a consequentialist account of the reasons agents have to regard themselves in certain ways. When an agent does something imprudent, recognizing his error and having certain emotional responses about it will likely make him less likely to make such errors in the future. And these practical reasons I mention are straightforwardly self-regarding, and therefore, non-moral. Again, this is intuitive. It would be bizarre to claim that he has moral reasons for thinking his actions foolish, or for being embarrassed, especially when we already grant that his action was not immoral in the first place (as I have already argued).

3. Now let's consider the reactive attitudes on the part of others in response to an agent's failure to appropriately respond to his other-regarding reasons, in other words, his having done something immoral. When an agent does something immoral, others generally have reason to blame him, to judge his character disfavorably, perhaps to be angry with him and, in extreme cases, to hate him. The reasons others have to take up these attitudes are other-regarding. But the relevant "other" here is not the wrongdoer, but the individual(s) whom he wronged. On my view, a moral wrong always involves one agent wronging another individual. Blaming an agent who wrongs another individual often demonstrates proper regard for the individual who was wronged. It is a way of recognizing the individual's moral status and the wrong that has done to him/her/it. This account of our reasons for blaming wrongdoers is consistent with my view, but that on its own is not much of a reason to accept it. It should enjoy its own intuitive appeal. Does it?

Suppose Clyde rapes Claudia. In general, some of us will have reason to blame Clyde. What reason? That Clyde seriously morally wronged Claudia. To see the intuitive appeal of this account, consider Clarke, who fails to blame Clyde. Now it is true, as we have seen from Smith, that one does not always have decisive reason to blame a wrongdoer. In other words, when an agent does something immoral, it doesn't always make sense for someone else to blame him, even if the wrongness of the act generates a reason to blame him. This is because there can be competing reasons for not blaming him, e.g., reasons involving standing, the significance of the wrong, the wrongdoer's own reactive attitudes, etc. But in our present case, suppose none of these reasons apply to Clarke. Suppose Clarke is Claudia's brother, so he has standing, in Smith's sense. We also know the wrong is quite serious. And suppose that Clyde is not at all remorseful. Rule out any reasons that might compete with Clarke's reasons to blame Clyde, such that Clarke has most reason to blame Clyde. Now consider Clarke's failure to blame Clyde for raping his sister. There seems to be something deeply wrong with such a failure. But what exactly is wrong with it? I think it is that, in

failing to blame Clyde, Clarke fails to properly regard his sister. We think that Clarke must not really care about his sister all that much, not just as his sister, but as a person. (Remember: the only reason I made Clarke Claudia's brother was so he would have standing.) We could consider many more examples. When one fails to blame Hitler for all the horrors he caused, one fails to properly regard all the people who were killed, or whose lives were devastated, as a result of The Holocaust. We should care about The Holocaust, because we should care about the people whose lives were destroyed by it.

These cases raise an interesting question: Is it sometimes appropriate to blame someone for failing to blame a wrongdoer? The answer that is both consistent with my view and with intuition is: yes. I have claimed that it is wrong for Clarke to fail to blame Clyde, and I have also said that when an agent does something morally wrong (or fails to do that which he morally ought to do), it is, in general, appropriate to blame him. It follows from this that it can be appropriate for some to blame Clarke for failing to blame Clyde, since it is wrong for Clarke to fail to blame Clyde. It certainly isn't the case that everyone has decisive reason to blame Clarke for failing to blame Clyde. There are many reasons agents may have not to blame Clarke. My point is that in some cases, it makes sense to blame an agent who fails to blame another agent who did something morally wrong.

One thing to note here is that when Clarke fails to blame Clyde, he is less blameworthy than Clyde is for raping Claudia in the first place. In other words, rape is much more morally wrong than the failure to blame a rapist. Though we might think that in some cases, failing to blame can itself be very wrong. There might be cases where failing to blame has serious consequences. We might think Joe Paterno's failure to blame Jerry Sandusky had tragic consequences. Two brief points here. First, it isn't quite accurate to say that Paterno's failure to blame had tragic consequences. Rather, it was his failure to report Sandusky. He could have reported Sandusky without blaming him, or he could have failed to report him while blaming him. Second, even if we say that it was Paterno's failure to

blame that had tragic consequences, we are still inclined to blame Sandusky *more* than we blame Paterno. Still, it could be that even though Sandusky deserves more blame than Paterno, Paterno's failure to blame was very wrong.

This in turn raises another interesting question: How long can this chain of blaming for failing to blame go? Should Clint blame Clarke for failing to blame Clyde? I have claimed that it is wrong for Clarke not to blame Clyde, so, yes, under certain circumstances, it would make sense for Clint to blame Clarke for failing to blame Clyde. But again, Clint's failure to blame Clarke is less serious than Clarke's failure to blame Clyde, and, of course, much less serious still than Clyde's initial act. Now, should Claire blame Clint if he fails to blame Clarke? I'm not sure. Clint's failure to blame Clarke for failing to blame Clyde does not seem all that serious. The point is that the wrongness will decrease as we move down the chain. If *A* commits a seriously immoral act, *B* may have decisive moral reason to blame *A*. In that case, it is wrong for *B* not to blame *A*. This means that *C* may have decisive reason to blame *B*. But *B*'s failure to blame *A* is less wrong than *A*'s initial act, and *C*'s failure to blame *B* is less serious still. So by the time we get to, say, *G*, it may no longer make sense for *H* to blame *G* for failing to blame *F*.

4. Finally, let's consider the reasons others have to take up certain reactive attitudes in response to another agent's failing to properly respond to his self-regarding reasons – that is, the reasons they have to take up certain attitudes in response to another's imprudent actions. I have already considered several such cases, e.g., the story of Jason who was tricked by an email scam. I have argued that his sending the money, though imprudent, is not necessarily immoral (depending on other details of the case such as whether he has any dependents). The main point I want to make here is that if we have reasons to make a judgment about Jason or his action, these reasons are not moral reasons. For example, it may make sense to have the attitude that Jason is a fool, or at least that Jason acted foolishly (he did act foolishly after all). But if we do have such reasons, these are not

moral reasons. To see this, suppose Lindsay, Jason's acquaintance, has decisive reason to judge that Jason is a fool, and yet does not. That is not immoral. It would be quite odd for someone to blame Lindsay for failing to make that judgment, even if she all things considered ought to. This is similar to someone who does something imprudent, something she all things considered ought not do, for the benefit of someone else, for example, by forgoing some benefit for the lesser benefit of another. Some actions of this kind are foolish and are such that all things considered, one ought not do them. But we certainly don't think there is anything immoral about them. We might think that all things considered, something has gone awry in the way Lindsay has formed her reactive attitudes regarding Jason, that, all things considered, she should judge that he is a fool. But it could very well be the case that there is in fact something morally praiseworthy about not making a harsh judgment of Jason, even if such a judgment is warranted. This is very different from the case I discussed above where Clarke fails to blame Clyde for raping Claudia. There, we saw that there is something morally wrong with Clarke's failure to blame Clyde. But it isn't morally wrong for Lindsay to refrain from judging Jason a fool for being taken in by an email scam.

Though we don't have moral reasons to judge that Jason is a fool, we may have moral reasons to take up other reactive attitudes in response to Jason's action. For example, we might have reason to take pity on Jason, even if he the cause of his own harm. I have already said that we have reasons to regard other individuals in certain ways. This, in some cases, may involve caring about their welfare. Jason has suffered some harm, and we might have moral reasons to care about that. It seems to me that I have reason to care that a person got hit by a truck, even if I discover that it could have been avoided if she had looked both ways before crossing the street.

I began this chapter by attempting to motivate the view that morality has only to do with how we treat others, and the attitudes we have about them. In this section, I've discussed reactive attitudes in detail. We have seen that not only is my view consistent with fairly strong intuitions

about what kinds of reactive attitudes we have reason to have in various cases, these intuitions lend additional support to the view I am defending, since, as I have shown in this section, alternatives are often very counter-intuitive indeed. In the following section, I will discuss several other considerations that help to motivate, or lend plausibility or further confirmation, to this view.

4. Miscellaneous Considerations

I am now going to consider several miscellaneous considerations that provide support for the claim that morality has to do with how we treat individuals other than ourselves. I will begin with Jesus' famous dictate, "Do to others as you would have them do to you" (*New International Version*, Luke 6:31) Many people take this not only to be sound moral teaching, but to be the sum of morality. Now, some may quibble with the idea that we ought to treat others how we would want to be treated, since people desire different things. If I love fishing, I would like it very much if a friend surprised me with a fishing trip. But it would be silly for me to do this for my friend who hates fishing. Set such concerns to the side for the moment. The point here is the object of our moral behavior. It is not ourselves. It is others. Jesus goes on to say,

If you love those who love you, what credit is that to you? Even sinners love those who love them. And if you do good to those who are good to you, what credit is that to you? Even sinners do that. And if you lend to those from whom you expect repayment, what credit is that to you? Even sinners lend to sinners, expecting to be repaid in full. (Luke 6:32-34)

The point he is making, I think, is that there is no moral worth in treating others well in hopes of receiving some benefit from them. This is why it is important to say that moral reasons are reasons we have to treat others well *for their sakes*. Lending someone money is a way of treating someone. But it is not a moral reason when the reason for lending the money is to earn some

interest, since you would be treating her in a certain way *for your own sake*. And in fact, we would not praise a banker who only went into business to make money by collecting interest on loans, even if he was not harming anyone, even if he was actually benefitting others by providing loans. Why? Because the reason he is loaning the money is to secure benefits for himself.¹⁴

Similarly, Jesus elsewhere instructs us to do our good deeds in secret rather than announcing them to the whole world (Matt. 6:2-3). There seems to be something to this prescription. But why? Again, it seems like the best explanation is that there is no moral merit in doing good deeds simply because we want people to know about them so they will treat us well.¹⁵ And notice also that treating people well in order that they treat you well certainly *is* a reason to treat people well. If it is a reason to treat people well, and yet there is no moral merit in treating people well *so they will treat you well*, it would seem that that is not a moral reason, that is, it is a non-moral reason¹⁶.

Although perhaps there is an alternative explanation for why we think we think it is morally better to keep our good deeds private. Maybe it is that people find it annoying when other people brag about their good deeds. Maybe it is not even in the interest of the good-deed-doer to broadcast her good deeds, since she does not want people to be annoyed with her. But this is not a good explanation. While it is true that sometimes boasting about one's good deeds backfires, it certainly

¹⁴ I do not claim that all bankers are purely selfish. Banks provide valuable services to many, including loans. If a person becomes a banker in part because he believes he can help people by providing them loans, that might be morally praiseworthy. But there is nothing morally praiseworthy about loaning money for purely self-interested reasons, even if the banker happens improves the lives of others by doing so.

¹⁵ One potentially embarrassing fact here is that Jesus says that if we perform our good deeds in secret, God will reward us in heaven. So perhaps Jesus doesn't think that there is no moral merit in acts we perform to secure some benefit for ourselves. But we need not worry about this. I am not making an appeal to authority. So we can divorce Jesus' directive to do our good deeds in secret from the promise of rewards in heaven. If we think there is something to the claim that it is morally better to do our good deeds in secret from a secular point of view, the reason we think there is something to it likely has nothing to do with rewards in heaven. So there can be alternative explanations as to why it is morally better to do one's good deeds in secret.

¹⁶ Similarly, we don't think it is morally meritorious when a politician volunteers at a soup kitchen when there is good reason to think he is only doing it to increase his popularity and get elected.

does not always backfire. It very often has its intended consequences, which is to gain the admiration and respect of one's peers. This is especially true if one can do it in an especially casual or graceful way, where it does not appear as though one is bragging. So if it seems that there is something less morally meritorious about doing good deeds in order to brag about them later to earn the respect and admiration of others, we should ask why that is so. The best explanation seems to be that the scope of morality is limited to our treatment of others, and therefore only reasons for treating others in certain ways for their own sake can properly be considered moral reasons.

Another consideration that seems to count in favor of this distinction is that many people recognize a distinction between the moral and the prudential. This distinction is assumed and referred to all the time, though the distinction is not often spelled out very carefully. So what exactly is this distinction? Well, the prudential is straightforwardly about the interests of the agent who is doing the action. To behave prudentially is for an agent to behave in a way that is good for her, or promotes her welfare. So what is left for the moral? Everyone else. To behave morally is for an agent to behave in a way that is good for others, or promotes their welfare.

This again does not prove that moral reasons have to do with our treatment of others while non-moral reasons have to do with how we treat ourselves. The distinction could be spelled out in different ways. While it might be true that my prudential reasons are limited to my own interests, perhaps morality considers *everyone's* interests, including my own. That could be right, but it does not seem to track the distinction people seem to be reaching for when they appeal to the moral/prudential distinction. It seems to me that the distinction people are going for is sharper.

My view also explains a long-standing misgiving I have had about utilitarianism. The complaint that utilitarianism is too demanding is of course quite well known. According to utilitarianism, we are very often morally required to satisfy the interests of others at the expense of our own. If I can raise the total amount of utility in the world by performing some action, but I

would be made much worse off by performing the action, morality, according to utilitarianism, requires that I do that action (assuming there are currently no other means of raising to the total amount of utility available to me). But there is an inverse objection to utilitarianism. It is that whenever I can raise the total amount of utility in the world, and I will be made better off in the process, I am morally required to do so. This might not seem like an objection on its face. If I can raise the total amount of utility by performing some action, and I will not even be harmed, in fact I will actually *benefit* from performing that action, then it might seem obvious that I am morally required to do so. But there are cases where we can sacrifice our own happiness for the happiness of others, and this will result in less utility in the world, and utilitarianism entails that this is morally impermissible. Suppose I have some possession that is of great value to me, a first edition of Ross's *The Right and the Good*. Owning this book gives me a lot of pleasure. Now I consider giving my book to a friend as a gift, to show her how much I value her friendship. I know that she will appreciate it, and that she will enjoy owning the book. But I also know that she will never enjoy owning the book as much as I do. I know that if I give her the book, the total amount of utility in the world will actually decrease. Under these circumstances, utilitarianism entails that it would be morally wrong for me to give her the book. But that just seems wrong. Of course it would be morally permissible to give her the book. Perhaps all things considered I ought to just find a different gift for her. It might not be the smartest decision I could make. But it is surely not morally wrong for me to give her the book.

Once again, my view is consistent with the clear intuition that it would not be morally wrong for me to give her the book, and it explains why it would not be wrong. It would not be wrong because I have moral reasons to give her the book, and non-moral reasons to keep the book for myself. Since non-moral reasons can help determine the moral status of an action, my non-moral reasons probably make it the case that I am not morally required to give her the book, but it would

also be morally permissible to give her the book, since I have moral reasons to do so¹⁷. Some might even consider the act *especially* praiseworthy or supererogatory, since I'm giving up a precious possession that I am not morally required to give up.

Two brief related points on supererogation. Suppose I am considering performing some action, and it would be morally better to perform the action than not to perform it, but it would be permissible not to perform it since it would be too costly for me. This is a case of supererogation. If there are actual instances of such cases, as there seem to be, it seems somewhat odd that though it would be morally better to perform the action, I am not morally required to perform it. My claim that self-regarding reasons are non-moral and other-regarding reasons are moral explains how this can be the case. My moral reasons count in favor of performing the action, but my non-moral self-regarding reasons count in favor of not performing the action, and sometimes justify not performing the action.

The other point about supererogation is that when people talk about supererogation, the examples they give never involve an agent's exceptionally good treatment *of herself*. There are many acts that we would consider supererogatory were the agent to perform them for the sake of another. But these same acts would never be considered supererogatory if done for the sake of the agent herself. We can consider again the example of giving away my expensive coat. If I were to give my new coat away, this would be a very kind and generous act – one that many would consider to be above and beyond the call of duty, that is, supererogatory¹⁸. But no one would consider it supererogatory to keep the coat for myself, even if I would not be able to acquire a nice coat without a generous act by me toward me.

¹⁷ Again, set aside concerns about what else I could do with the book, like sell it and donate the money to Oxfam. For example, assume that the market value of the book is quite low.

¹⁸ That is, at least under certain circumstances. If I have one old, worn coat, it could very well be permissible to replace it with a newer, nicer coat, even if I don't strictly speaking need it for warmth. If I already have five nice coats, perhaps then I really should buy a new coat for someone in need.

Once again, my view explains why this so. I have moral other-regarding reasons to buy a coat for another person, and while I certainly have reasons to buy myself a coat, these reasons are not moral reasons. This explains why it is perfectly permissible for me to buy myself a coat, and why it may even be the case that I *ought* to buy myself a coat, even though there would not be anything particularly morally good or praiseworthy in doing so. Again, it is because the reasons for doing so are non-moral.

Finally, consider cases where someone does something stupid or imprudent for the sake of someone else. Imagine that Mother's Day is around the corner, and Greta is considering what she should get for her mother. She knows her mother has fantasized about owning a sporty Mini Cooper. Greta rightly loves her mother very much and she decides to buy her mother the car for Mother's Day, even though she cannot afford it. She will have to take out a high interest loan that will take a long time to pay off. Greta has certainly made herself considerably worse off by buying this car for her mother—the harm she will incur outweighs considerably the benefit her mother will get from the car. Greta should not have bought the car for her mother. It was a stupid thing to do. But setting aside concerns about what else Greta might have done with that money (donating it to charity), did Greta do something morally wrong? Were the reasons for not buying the car for her mother primarily moral? Once again, common sense says that what Greta did was not morally wrong, and her reasons for not buying the car are not moral reasons. What we would normally say about Greta is, "Her heart was in the right place," or "Her heart was bigger than her brain."¹⁹ We would only say that someone's heart is bigger than her brain when her concern for others moves her to do stupid or imprudent things, or things that she all-things-considered should not do. And we don't mean to be making a negative moral evaluation of the person or her actions. Quite the

¹⁹ This is what Clark Griswold says of his moronic brother-in-law in "National Lampoon's Christmas Vacation."

opposite. We mean that the agent is morally admirable – the action was morally good – even if it was less than fully rational.

But this is hard to explain if moral reasons can be self-regarding. If Greta's reasons for not buying the car are moral reasons, and we agree that she has more reason not to buy the car than to buy the car, then it would seem that Greta has most moral reason to refrain from buying the car. That suggests that she is morally prohibited from purchasing the car and that doing so would be morally wrong. And that would be a strange assessment of Greta's action. But if Greta's reasons for buying the car are moral reasons, and her reasons for not buying it are not moral, then we can claim that, although she has most reason to refrain from buying the car, doing so was not morally wrong.

Yet another advantage of the view that self-regarding reasons are non-moral is that it helps to solve a puzzle faced by those who endorse certain features of common sense morality. This puzzle is expressed most forcefully and explicitly by Michael Slote. Slote notes, "Ordinary morality (unlike act-consequentialism) treats our obligations to others as dependent on how near they stand to us in relation of affection or special commitment" (184). So we might be inclined toward a view that says that the closer the relation in which a person stands to an individual, the more pressing the agent's obligation is to that individual, or the stronger the agent's reasons are for treating that individual in a given way, or the more the agent must consider that individual's interests. But it is also part of common sense morality that "it is quite permissible to sacrifice one's own greater benefit to the lesser benefit of another" (180). This is part of a more general phenomenon that has come to be known as the self-other asymmetry, which is an asymmetry "regarding what an agent is permitted to do to himself and what he is permitted to do to others" (181). Another example that Slote mentions is that it may not be permissible for an agent to cut up a person in order to distribute his organs to those who need them, but it would be perfectly permissible for him to do this to himself

(or, rather, to allow that this be done to him).²⁰ While it may be foolish or crazy for an agent to sacrifice a benefit for himself to the lesser benefit of someone else, it is surely not morally wrong, according to ordinary morality. I of course would go even further and say that it is *never* morally wrong for a person to make herself worse off, so long as she does not make (or intend to make) anyone else worse off in the process, and especially if she makes others better off in the process. But these two features of ordinary morality are in apparent conflict. Slote writes:

On the one hand, [common sense morality] encourages the idea that strength of obligation weakens as one gets further from the agent, but, on the other hand and in seeming opposition to the first idea, it assumes that there is no moral obligation whatever (except indirectly) for the agent to benefit himself or concern himself with his own projects. (185)

So this is a puzzle: On the one hand, we want to say that our moral obligations to another person are stronger the closer they are to us, while at the same time, we have no moral obligations to our own selves, and to whom are we closer than our own selves? This is especially salient here since my view does entail that we have special obligations to those with whom we have a special relationship, and, by and large, I am attempting to preserve common sense morality as much as possible. That is, I take common sense morality seriously.

But saying that an agent's reasons for treating herself in certain ways for her own sake are not moral reasons, but are reasons nonetheless, helps to solve this puzzle. Although our relational reasons are often stronger than our non-relational other-regarding reasons, and while it is of course true that we stand in quite a special and unique relationship to ourselves, it is not morally wrong, for example, for an agent to sacrifice a benefit to herself to the lesser benefit of another (even a

²⁰ Slote later explains why he thinks consent does not account for the difference in permissibility. I agree that consent does not account for the difference, but I will consider cases like these in detail in later chapters.

complete stranger), since one's reasons for securing the benefit for herself are not moral reasons. Of course, on my view an agent still has reasons to treat herself in certain ways, to try to promote her own welfare, etc. And in some cases, these reasons might even be weightier than the reasons she has to promote the welfare of a stranger. Hence our conviction that sacrificing a benefit for oneself to the lesser benefit of another is (at least in some cases) foolish or crazy. It might be that the agent, all things considered, ought not do this. But it would not be morally wrong.

Now, this doesn't answer the question: Why are self-regarding reasons non-moral? As I have suggested, I don't think there is an answer to this question. On my view, it's just a brute normative fact. But the point here is that it is theoretically advantageous to claim that self-regarding reasons are non-moral, because it makes sense of the two commonsense notions that while our moral obligations are more pressing the closer the relevant individual is to us, we have no moral obligations to treat ourselves in any way for our own sakes.

The claim for which I am currently attempting to provide support is that moral reasons are only those reasons for treating others in certain ways, for their own sakes. Moral reasons are never reasons an agent has to treat herself in certain ways, for her own sake. What I have tried to show in this chapter is that this distinction is consistent with and, further, often explains a lot of intuitions and assumptions that we take on, perhaps not consciously or explicitly, when thinking and talking about morality and practical rationality. Alternative explanations may be possible in each case, but the view I am offering seems to explain all of these cases in an intuitively plausible and elegant way.

In the next chapter, I will consider objections to these claims.

CHAPTER 3

Moral Reasons as Other-regarding Reasons: Objections

In the previous chapter, I argued that moral reasons are reasons to treat others in certain ways and have attitudes regarding them for their own sakes. A catch phrase for this view is: ‘Morality is about how we treat others.’²¹ In this chapter, I will consider objections to this view.

1. Objections to the Test

I have claimed that one way to test the moral status of an action is to ask what sort of reactive attitudes would be appropriate in response to it. In general, when someone commits an act that is morally wrong, it is appropriate for her to feel guilty, and it is appropriate for others to blame her or feel righteous indignation. But these attitudes are not appropriate when the agent has merely done something foolish or imprudent. Likewise, when someone does something that is morally right, it is appropriate for others to praise her, and for her to have a clear conscience and a sense of moral pride (or whatever the opposite of guilt is). Because we do not blame people when they perform actions that only harm themselves, we should not think of such actions as morally wrong. Likewise, because we do not praise people for performing actions that only benefit themselves, we should not think of such actions as morally right or as having any moral merit. I will here consider two objections to this test.

The first objection is the claim that actions that only benefit the agent performing them can be morally right, but they do not necessarily warrant praise. This is because people are already *so* inclined to benefit themselves that we do not need to encourage this sort of behavior. So one’s

²¹ But it is just that: a catch phrase. I will qualify this in the final chapter.

reasons to benefit oneself may be moral reasons, but there are other moral reasons that too often get ignored, so we should be encouraging people to act on those reasons.

My response to this objection is that, as we saw in the previous chapter, praise is not always, or even often, about encouraging behavior. Praise, as I am using the term, is a reactive *attitude*. Just as blame, even active blame, need not ever be expressed, neither does praise. It can be entirely internal. I've never discussed my opinions about the Dalai Lama with anyone, and yet it could still be true that I praise him or admire him.²² But we don't even privately praise people who do actions that only benefit themselves. Our estimation of their moral character does not go up. It simply seems true that there is nothing morally meritorious about such actions.

Maybe the objection is not so much about encouraging behavior, but rather that there is such a strong natural tendency to be highly responsive to our self-regarding reasons, or to benefit ourselves, that there is nothing extraordinary about it, so we are often not inclined to praise such behavior, even if it does have moral merit.

The problem with this objection is that there are ways in which we are responsive to other-regarding reasons, or actions we do for the sake of others, that also come quite naturally, but also warrant praise. If Meghan sees a child drowning in a shallow pond, she might rush in to save him without a second thought. Such behavior may come quite naturally, and yet I still think it would be appropriate for others to praise Meghan. It is often quite natural for parents to promote the welfare of their children, but we still praise good parents.

²² Admittedly, 'praise' can have the connotation of being something that we do, that is, something we express externally. This seems to be true of 'praise' in a way that is not true of 'blame.' It seems more natural to say I can blame someone without ever opening my mouth than that I can praise someone without opening my mouth. 'Admiration' certainly seems like it can be an attitude that is never expressed, but 'admiration' does not seem as morally tinged as praise, as in the case of a secret-admirer (though I suppose the concept of non-moral praise is coherent as well). But since praise seems to be the natural opposite of blame, I will continue to use praise in such a way that it can be a private, unexpressed attitude.

Praise, as I am using the concept, does not have to be some wildly enthusiastic judgment about a person or a person's actions. It can be a private, sober judgment. Simply forming the judgment that Meghan did the right thing when she saved that child drowning in the pond or that James is a good father can be forms of praise.

The second objection is simply the converse of the first. It might be true that there are moral self-regarding reasons, but we do not blame people when their actions only harm themselves, since they have already suffered enough, and do not need to be discouraged from that behavior in the future.

My first response here will be familiar. Blame does not have to be publicly expressed, and it need not have anything to do with attempting to change behavior. It can simply be a private judgment. But we don't privately, or passively, blame people when they fail to appropriately respond to their self-regarding reasons. It simply seems true that they have not done anything morally wrong or blameworthy. Perhaps there is such a thing as non-moral blame, which would be making a negative assessment of some aspect of a person, but one that has nothing to do with her moral character, i.e., judging that someone is stupid or foolish or imprudent. We do often make these kinds of judgments of people who fail to appropriately respond to self-regarding reasons. But 'blame' (*simpliciter*) has clear moral connotations so I don't think it would be appropriate to classify these judgments as instances of blame.

Furthermore, it is appropriate to blame people when they fail to properly respond to other-regarding reasons, even when they themselves are the only ones harmed. If a jewel thief is scaling a building and falls and breaks his leg and is then arrested, it is still appropriate to blame him. So the fact that a person who fails to appropriately respond to his self-regarding reasons often bears the brunt of the ensuing harm does not explain why it is not appropriate to blame people who fail to respond appropriately to their self-regarding reasons.

2. Counter-examples

Some objections to the claim that morality has only to do with how we treat others for their own sake may come in the form of counter-examples. The cases I have been considering to motivate this claim involve an agent harming herself in some way, and I have been claiming that it is not appropriate to blame the agent in these cases. There may be some who agree with the assessments of the particular cases, but still think there are ways an agent can treat herself, for her own sake, that would be morally wrong, or might warrant blame, or might think the agent's behavior is personally discrediting. I will now consider several possible examples of such cases.

2.1 Self-regarding Racism

The first is what we might call "self-regarding racism." Consider a racist who sincerely believes his own race is inferior, and this belief motivates him to harm himself. We could imagine the blind racist from *Chappelle's Show*, who is aggressively racist against blacks, while he himself is black, but he is unaware that he is black. Suppose it is eventually revealed to him that he is in fact black, but his racism remains. Since he has a standing belief that black people deserve to be harmed, he harms himself. Perhaps he commits suicide.

Many would probably be inclined to think this blind racist is morally in the wrong and he and/or his actions and/or his attitudes warrant blame or indignation. It certainly seems clear that he was wrong for being a racist before he knew that he himself was black. Why think the wrongness disappears upon his discovery? If his harming himself is morally objectionable, then my view is false, since he is only harming himself, and on my view, it is never morally wrong to treat yourself in any way.

This difficulty can be overcome when we separate out the act from the motivation and consider each on its own. As we have seen, it is often appropriate to evaluate a person's attitudes

and beliefs. In the previous chapter, I gave both consequentialist and non-consequentialist explanations for why this is so. While either is consistent with my view as a whole, I endorse the view that attitudes can be wrong in themselves, that is, there are non-consequentialist reasons to regard certain individuals in certain ways. But my response here will be successful even if there are only consequentialist reasons to not have racist attitudes. For example, one might think one should not have racist attitudes because having such attitudes makes it likely that one would eventually mistreat others.

If it is true that there are moral reasons to not have racist attitudes, then I can consistently maintain that there is something morally problematic about the situation as a whole. The blind racist is morally in the wrong for his racist beliefs, but his act of self-harm is not morally wrong (though, of course, he has reasons to refrain from self-harm, and he should probably not harm himself). This is consistent with my view, since his racist attitudes are clearly other-regarding, especially since, before his discovery, he does not even know he is black²³. All of his ire is directed outward. When he does discover that he is black, and is therefore motivated to harm himself, his act of self-harm is not morally wrong, though he does have reason not to harm himself. So it is true that there is something morally problematic about the situation as a whole, namely, his racist attitudes. It need not be the case that his act of self-harm is morally wrong in order for it to be true that there is something morally problematic going on.

And this response is not *ad hoc*, since it accords with what intuitively seem to be the warranted reactive attitudes. The blind racist really is morally blameworthy for his racist attitudes, precisely because these attitudes are other-regarding. When he discovers he is black and is motivated

²³ This is consistent with my view even if the racist is not blind and is racist against her own race, regardless of whether the reasons to not be racist are consequentialist or non-consequentialist. If she is racist against her own race, she is more likely to mistreat those of her own race. She may also fail to respect them or properly regard them in some other way.

to harm himself, I personally feel pity for him, not blame. At the very least, it is much less clear that blame or righteous indignation is warranted for his act of self-harm, so we should not rely on such shaky intuitions to undermine the hypothesis that morality only involves our treatment of others, when there are so many other cases that confirm this.

One possible objection to this solution is that because the blind racist's motivation (i.e., his racist attitudes) is morally objectionable, then his act was morally wrong as well. The success of this objection depends on the truth of what Liezl van Zyl calls "strong motivism," which is the view that "a good motive *always* serves to make an act right, whereas a bad (or insufficiently good) motive *always* makes it wrong" (2010, 407, emphasis added). This view is to be distinguished from "weak motivism, which is the view that motives only sometimes help to determine the deontic status of an action, as well as from "anti-motivism," which is the view that motives *never* play a role in determining the deontic status of an action. As van Zyl points out, strong motivism is not a popular view, and for good reason. Suppose I see a child drowning in a shallow pond but I am disinclined to save him, but then it occurs to me that if I do, I might end up on television or get a reward. Could the fact that I would be acting from a bad (or insufficiently good) motive entail that it would not be right to save him? That it isn't the case that I ought to save him? That cannot be true. Of course I should save him, regardless of my motives. So that leaves weak motivism and anti-motivism, both of which are consistent with the view that although the blind racist acted from immoral motives, his act of self-harm was not morally wrong.

2.2 Self-Degradation

Another way in which one might think one can morally wrong oneself is by degrading oneself. The obvious putative examples here would be practices like pornography and prostitution, but there may be others. Another example of degradation might be a practice that has come to be

known as “dwarf tossing.” According to Wikipedia, “Dwarf tossing is a bar attraction in which dwarfs wearing special padded clothing or Velcro costumes are thrown onto mattresses or at Velcro-coated walls.” Apparently, the goal is to throw the dwarf higher than anyone else. Of course I do not deny that pornography and prostitution can be harmful to various parties involved, and it is consistent with my view that participation in these activities is often or perhaps even always morally wrong. And there does seem to be something deeply troubling about dwarf tossing. But the key question here is, assuming that participation in these activities is (at least sometimes) morally wrong, whether it is wrong *because* it is degrading and, even more importantly, whether the prostitute or the pornographic actor or the dwarf who allows himself to be tossed degrades himself or herself and thereby wrongs himself or herself. Is it plausible that self-degradation is both possible and morally wrong?

To answer these questions, we must first know what degradation is. Before discussing this, I simply want to mention the possibility that degradation may be one of those terms that resist conceptual analysis. It might be that, in fact, there is no such thing as degradation. If that is the case, then one cannot morally wrong oneself by degrading oneself, and so much the better for my view. While the view that there is no such thing as degradation does not strike me as implausible, I will not avail myself of this strategy. Instead I will focus on what I take to be the most plausible analysis of the term and see what follows.

The concept of degradation has not received all that much attention from analytic philosophers. It receives its most direct and careful attention perhaps in Judith M. Hill’s paper “Pornography and Degradation.” According to Hill, it is a necessary condition of degradation that one be “treated as though one were a means only, as though one were not an end in herself, as though one were something less than a person” (1987, 40). But that is not sufficient, according to Hill. It is also necessary that this treatment be public. Referring to prisoners in Auschwitz, she writes

that shaving the heads of the prisoners was degrading “because it marked one as a member of the class that was being treated as sub-human.” She continues,

Forced marches for the prisoners doing housekeeping chores were degrading because their sole purpose was to exhibit—for the benefit of the prisoners and the guards, at least—the complete submissiveness and obedience of the prisoners and the complete control of the guards...[A]n agent who treats his victim as something less than a person in public places, for the whole world to observe, demonstrates a conviction that her worthlessness is so extreme that all the world can be counted on to regard him as justified in treating her accordingly. (41)

If ‘degradation’ refers to anything, that seems about right. One thing to note is that this case, and in fact all of the cases Hill discusses, are cases where one person is degrading another person. This of course does not entail that one cannot wrongfully degrade oneself, but once again, this suggests that paradigm cases of wrongful degradation are those that involve one person degrading another.

The relevant question here is not whether a john wrongs a prostitute by degrading her, or whether an average-sized person wrongfully degrades a dwarf by dwarf-tossing. The relevant question is whether the prostitute wrongs *herself* by degrading herself by being a prostitute, or whether a dwarf wrongfully degrades *himself* by allowing himself to be treated in this way.

For the moment, I will simply assume that dwarf-tossing is in fact degrading, and even that a dwarf degrades himself by allowing himself to be tossed. To use Hill’s language, the dwarf who allows himself to be tossed allows himself to be treated as something less than a person in a public place and “demonstrates a conviction that [his] worthlessness is so extreme that all the world can be counted on to regard [the tosser] as justified in treating [him] accordingly.” The question here is whether the dwarf who allows himself to be tossed wrongs himself by degrading himself.

One strategy we might take in answering this question is to ask what is wrong with degradation in general. A necessary condition of self-degradation's being a way in which one can wrong oneself is the truth of two propositions:

- (1) The wrongness of the degradation is not explained by its being harmful.
- (2) The wrongness is not fully mitigated if the degradation is consensual.

The truth of (1) is a necessary condition of self-degradation being wrong because, as I have already argued at some length, it is implausible that one wrongs himself by harming himself. So if degradation is morally wrong just because it is harmful (in other words, because it is a way in which one can harm oneself), one does not morally wrong oneself by degrading oneself, since it is not morally wrong to harm oneself.

Why the truth of (2) is a necessary condition of the wrongness of self-degradation should be fairly obvious: Obviously if I degrade myself, or allow myself to be degraded, I am consenting to it. (That is assuming I am not being coerced into degrading myself, but that would be problematic for a different set of reasons.) So if it is not wrong to degrade someone when you have his consent, then it is not wrong to degrade yourself when you consent to it.

However, deciding whether (1) and (2) are true is no easy task. (1) is false if the following proposition is true: (3) It is impossible to wrong someone without harming her. Whether (3) is true is itself a difficult question. As I stated in the previous chapter, I myself do not endorse (3), since I believe you can wrong a person by failing to properly regard her, even if your attitude has no effect on her. Still, I do not think it would be unreasonable to endorse (3), and many philosophers *do* endorse (3), and my view is not incompatible with (3). And if (3) is true, then it is not the case that one can wrong oneself by degrading oneself, and so much the better for my view.

Determining the truth of (2) is also rather difficult. (2) is false if the following proposition is true: (4) It is impossible to wrong an individual by performing an action if you have that individual's

informed consent that you perform the action (in the common law, this principle is expressed by the phrase *volenti non fit injuria*, sometimes simply referred to as The Volenti Maxim) . Unfortunately, I have already argued that (4) is not plausible. Consider again the scenario where I swim out to a drowning man in order to save him, and he tells me not to save him. He explains that he does not deserve to live because he has lustful thoughts about women. It seems that it would still be wrong to let him drown, and that I would be wronging him, even if I have his consent to let him drown.

Of course, the truth of (4) is not the only way (2) could be false. Degradation could be a special kind of wrong, where the wrongness *is* fully mitigated when it is consensual. Breaking promises seems to be this way. It is *prima facie* wrong to break a promise, but it could be that the wrongness is fully mitigated when you have the promisee's consent. This is assuming there are no other reasons to keep the promise, e.g., that it would harm the promisee. If you make a promise to someone, and the only reason to keep the promise is that you made the promise, and the person consents to your breaking the promise, then the wrongness of promise-breaking in this case might very well be fully mitigated. If she truly does not mind and nothing bad will happen as a result of your breaking the promise, it is fully permissible to break the promise. The same might be true of degradation. If you have someone's permission to degrade him, and there are no further reasons to refrain from degrading him, e.g., the degradation would not be harmful, then it very well might be permissible to degrade him.

So the proposition that self-degradation is sometimes wrong must jump the hurdles of both (1) and (2) in order for it to be true. I have not argued that it is unable to jump these hurdles, but I have tried to point out some of the difficulties in jumping these hurdles. But while jumping these hurdles is *necessary* for its truth, it is not *sufficient*. That is, (1) and (2) could both be true and it could still be false that self-degradation is sometimes wrong; (1) and (2) do not together entail that self-degradation can be wrong. And I argue that it is not wrong, even if (1) and (2) are true.

The strategy should by now be familiar. Imagine a case where a person degrades himself, and no one else is affected by the degradation. In such a case, it does not seem appropriate for anyone to blame that person. That person is not worthy of blame. Rather, pity seems to be the appropriate response. It indeed seems wrong to treat another person as though he is less than a person, perhaps even when you have his consent. The person doing the degrading might very well be open to blame. But when a person believes he is worthless and treats himself as though he is worthless, it is difficult to see how he is open to blame. His belief may be false, and his actions based on the belief imprudent, but not blameworthy.

Even if it is wrong to degrade someone when you have his consent, it is likely that there are often features of cases involving one person degrading another that can explain the wrongness that self-degradation will lack. For example, it can often be cruel to degrade someone, even if he consents. I suggest that this is often one of the features of the degradation that we are responding to when we have the intuition that it can be wrong to degrade someone even with his consent. It seems cruel for an average-sized person to participate in dwarf-tossing, even when the little person consents to it. In some cases, this might be in part because he is exploiting the little person's low self-worth. Another feature often present in cases involving one person degrading another that helps explain the wrongness is that the one doing the degrading often fails to properly regard the person whom he is degrading, and I have already argued that it is plausible that it is wrong to fail to properly regard other individuals. But I have also argued that although there is a proper way for an agent to regard himself, he is not morally wrong when he fails to properly regard himself. That the person doing the degrading often fails to properly regard the person he is degrading may also be helping to generate the intuition that degrading another person is morally wrong, even when the degrader has the degraded's consent.

If there are ways to degrade someone that are not cruel and that do not show a lack of proper regard, it is plausible that such degradation would not be morally wrong. Consider a man who visits a dominatrix.²⁴ For simplicity's sake, suppose no sex is involved – it is just the domination and humiliation. These scenarios can plausibly be described as instances of degradation (in fact, it is often explicitly the man's goal to be subjected to degrading behavior), but it is not clear that the dominatrix is doing anything wrong in degrading her client. I suggest that this is because the degradation is neither cruel nor expresses a lack of proper regard for the man (and it may lack other wrong-making features often present in degradation). Rather than expressing cruelty, the dominatrix can plausibly be described as providing a service to her client. Moreover, there is no reason to think that the dominatrix could not have all the proper attitudes toward her client that all of us ought to have regarding other persons.²⁵

Given what I have said about degradation so far, could it still be wrong for a little person to participate in dwarf-tossing? And if so, would this threaten my claim that moral reasons are not self-regarding? The answers to these questions are 'yes', and 'no', respectively. If it can be wrong for a

²⁴ It is true that sessions with a dominatrix are often not public, so according to Hill they may not count as genuine degradation. But Hill might simply be wrong that publicity is a necessary condition of degradation, or we can simply imagine a dominatrix scenario that *is* public. Nothing I say hangs on whether the degradation is public or private.

²⁵ I point out once again that examples in this context are often imperfect. I can think of two ways a reader might be unhappy with this example.

1. Some might not be persuaded that the scenario I have described is a case of real degradation. This could perhaps be because the man is paying the dominatrix, so in fact, he is the one who is in control. But it isn't clear why this fact would make it true that it is not a genuine instance of degradation. Moreover, if one thought that it is impossible for a person to degrade himself, or voluntarily allow himself to be degraded, then one would already be committed to the claim that a person cannot wrong himself by degrading himself or voluntarily allowing himself to be degraded, since it is conceptually impossible to do so.

2. On the opposite end, some might still think visiting a dominatrix is morally wrong. However, I suspect that most philosophers would not think this, at least as long as we stipulate that the man wants it or enjoys it, that the dominatrix regards him in the proper way, etc. But many non-philosophers might still believe there is something morally wrong with visiting a dominatrix. This concern will get sorted out below when I discuss the differences between the ways philosophers and non-philosophers approach moral questions.

dwarf to participate in dwarf-tossing, and for people to degrade themselves and allow themselves to be degraded in general, the wrongness can be explained by other-regarding reasons. For it seems to me that the best explanation of the wrongness on the part of the dwarf who participates in the practice has to do with *other* dwarfs. When a dwarf participates in dwarf-tossing, it contributes to and perpetuates an attitude in the public that dwarfs are appropriate objects of amusement and targets of ridicule. In this way, a dwarf who participates in dwarf-tossing harms other dwarfs, and other dwarfs would be justified in resenting him for his complicity.

Contrast this with a world that contains only one dwarf. In this world, we are certain that this dwarfism is completely anomalous – the result of a unique set of circumstances - and that no instances of dwarfism could ever occur again. Suppose this dwarf degraded himself by allowing himself to be made an object of amusement and ridicule. It does not seem like it would be morally wrong for the dwarf to behave in this way. And the view that morality is basically other-regarding offers a clear and straightforward explanation as to why it might be wrong in our world for a dwarf to participate in dwarf-tossing, but not wrong in the world that contains only one dwarf. In our world, the dwarf's behavior has implications for other dwarfs, but it does not in the one-dwarf world.

The view that it is wrong to degrade yourself when that degradation would have implications for other individuals is consistent with the best explanations of the wrongness of pornography. Of course, when it comes to pornography, there are many players involved: The male actors, the female actresses, the producers, the directors, the distributors, the consumers, etc. We need not make the same moral assessment of all of these players. It may be wrong to act in porn, but not to view it. Or they may both be wrong, but wrong in different ways. Most of the philosophical treatment on the subject does not focus on whether there is anything morally wrong for a woman to allow herself to

be the subject of pornography, but that is the relevant question here.²⁶ For example, Hill focuses on cases that involve one person degrading another, but these have no bearing on whether it is morally wrong for a person to degrade herself.

The most plausible explanations of the wrongness of pornography do not focus on the individual woman who is depicted, but rather on women in general. For example, Lynne Tirrell argues that pornography “serves to reinforce patterns of oppression, exploitation, and discrimination.” She continues, “Derogatory words and images have the power that they do because they support and are supported by a host of other sorts of social practices, from discriminatory employment practices to the redefining of neighborhoods to unequal access to education and so on” (1999, 231). Catherine MacKinnon and Alisa L. Carse argue along similar lines (1989). All of them focus on the harm pornography causes to women in general. If it is wrong for an individual woman to allow herself to be depicted in pornographic works, it seems wrong because she is participating in a practice that makes all (or most, or many) women worse off. The moral reasons she has to refrain from this behavior are other-regarding.

In sum, it is not morally wrong to degrade yourself, or allow yourself to be degraded, when no one else will be affected by the degradation. Self-degradation seems to be morally wrong only when the degradation will affect others. In other words, self-degradation seems wrong only when one has other-regarding reasons to refrain from it. If that is right, then self-degradation is not a counter-example to the claim that moral reasons are exclusively other-regarding.

²⁶ We might also ask about the male actors, but given contingent features of our world having to do with the differences between how men are viewed and treated and women are viewed and treated, it makes more sense to focus on the women depicted in pornography.

2.3 “Seedy” Behavior

Many people seem to think that drug use, gambling, and certain kinds of sexual behavior are (sometimes) morally wrong, and these might be behaviors where only the agent who is doing them is affected, or is at least the person who is primarily affected. As Michael Huemer points out, one feature all of these kinds of behavior have in common is that they involve the agent giving into some kind of temptation. I will call these behaviors *seedy* behaviors²⁷.

Are seedy behaviors morally wrong? Do we have moral reasons to refrain from doing them? The first thing to note is that on my view, they often are morally wrong, or at least we have moral reasons to refrain from doing them. If I take certain drugs, I may become addicted. A drug-addicted person is likely to cause a lot of harm to others. He may become a thief to support his habit. He may cause distress to his loved ones. He may be unable to support those who rely on him. He may not be able to contribute to society. So my view is consistent with the claim that one often has moral reasons to refrain from certain recreational drug use. The same goes with sex and gambling. There are many ways in which these behaviors can affect people other than the agent who is doing them.

For seedy behaviors to truly be counter-examples to my view, it must be plausible that these behaviors can be morally wrong, or that an agent can have moral reasons to refrain from doing them, even in cases where the agent doing them is the only individual who could reasonably be predicted to be affected by them. But once we consider cases where it is stipulated that only the person doing the behaviors can be predicted to be affected by them, the moral wrongness seems to vanish, or, at a minimum, it is much less clear that the behavior is morally wrong.

²⁷ It is true that ‘seedy’ is often used as a thick term (more on thick moral concepts below). There is badness built into the term. To say that a certain behavior is seedy usually implies that it is bad in some way. Here, I am introducing it as a morally neutral technical term. It is an open question whether these behaviors are actually bad, or how they are bad.

I have already mentioned the challenge of considering behaviors where only the person doing the behaviors is affected, which is that in the real world, our actions are not performed in isolation, and almost always have some affect on other individuals. This seems to be especially true when it comes to seedy behaviors. Gambling, for example, necessarily involves other people. Even with slot machines, my winning a jackpot will be at the expense of several others who have thrown their money away, and if I win, there will be less for others to win in the future. So it is impossible to consider cases of gambling in isolation.

Sex often necessarily involves others, so it is easy to explain the wrongness of certain sexual acts by the way they affect others. The obvious exception is masturbation, but the view that masturbation is wrong in itself is so out of the mainstream that I'm not going to consider it.²⁸ Now, it is more plausible that consuming pornography is wrong, and we might think this is a solo activity as well. But of course pornography does indeed involve others and as I argued above, the best arguments for the moral wrongness of pornography are those that appeal to the ways in which it harms women in general, not the consumer or the individual performer.

While drug use almost always affects those other than the drug user, we can imagine another last-man type case. If I was the last sentient being on earth and I found a large cache of recreational drugs, it is hard to see how taking them would be morally wrong.

I have been claiming that seedy behaviors are not morally wrong in a basic way, or not morally wrong in themselves. Perhaps many academic philosophers can get on board with this claim, but “lay people” (those not regularly engaged with academic philosophy) may continue to think that seedy behaviors are morally wrong in a basic way. And this may be because lay people just mean something different by ‘moral’ than philosophers do.

²⁸ I won't consider it except to say this: If one thinks that masturbation is wrong because it is degrading, this concern has already been handled in the section above on self-degradation.

Michael Huemer suggests something along these lines in his article “Values and Morals.” He points out that ‘morality’ is often used in ethics as a thin evaluative term, that is, a term that is “purely evaluative, making no particular demands on the descriptive features of a thing” (2009, 114). “In ordinary language” Huemer notes, “such terms as ‘moral,’ ‘immoral,’ ‘morally wrong,’ and ‘duty’ function as thick evaluative terms, with substantial semantic content beyond that of such thin evaluations as ‘good,’ ‘bad,’ and ‘something we have a reason for doing’” (119). The semantic content here is, according to Huemer, closely associated with behaviors I have been calling seedy.

One way to see the difference between the way in which lay people use the term and the way in which philosophers use it is to consider what each takes to be paradigmatic cases of moral wrongness. As Huemer points out, in academic ethics, a paradigm case of moral wrongness “would be causing suffering to another being, without that being’s consent, and without the expectation of achieving any compensating value” (120). He goes on: “In popular discourse, however, ‘immoral’ is most often used to describe such things as female sexual promiscuity, sodomy, excessive drinking, recreational use of drugs, and gambling.” He points out that in the Bible, the word ‘immorality’ is almost always preceded by ‘sexual’. He notes,

In popular parlance, the phrase ‘a woman with loose morals’ does not refer to a woman who causes unnecessary pain to others, violates others’ rights, or fails to assist others in need when she could do so at minimal expense; it refers, rather, to a woman who is unusually receptive to sexual propositions. (120)

Huemer thinks there is “great error” in conventional morality (and I agree), but he suggests that perhaps lay people are just after a different concept altogether than what academic philosophers are after. He provisionally concludes,

It appears, then, that ‘immoral’ may function as a thick evaluative term: one with negative evaluative valence, with an associated attitude of righteous disapproval, and

with a descriptive profile involving a giving in to temptation of some kind, especially temptations posed by sensory pleasure. (121)

If Huemer is right about all of this, then I am in trouble. I have been offering a particular analysis of morality, and have relied on what I take to be common intuitions regarding the various cases I have presented—intuitions I suspect are held by both philosophers and non-philosophers alike. So I take my analysis to jibe with ordinary intuitions about morality, and I take the fact (if it is a fact) that my view jibes with ordinary intuitions to *support* my analysis. But if Huemer is right, then common intuitions about morality often *conflict* with my analysis. Perhaps more importantly, I'm not even giving an analysis of morality at all, at least as the term is used in common parlance, but of something else.

I do not deny that there are many people who believe that drug use, promiscuous sex, and gambling are morally wrong in themselves, while most academic ethicists do not. Huemer's explanation for the disparity is that the two groups are just employing different concepts. He writes that contemporary ethics and common notions about morality "are distant enough that it is not implausible to suggest that academic ethical theories are not, after all, theories of *morality*, as the term is popularly understood" (124). I offer an alternative explanation. This alternative explanation is that we are reaching for the same basic concept, but we just have different beliefs about it.

What concept a person is employing by using a certain term is a deep and complicated question. It is not always clear whether two people mean the same thing when using a certain term. Suppose Jane and John are having a debate about knowledge. John thinks that knowledge is justified true belief. Jane thinks that while those conditions are necessary for knowledge, they are not sufficient. Are they talking about the same thing? Probably, but it may be difficult to determine. Upon observing their discussion, we might determine that their respective uses of the term are intuitively "close enough" to count as their talking about the same thing. But if Jane were to say,

“You think knowledge is justified, true belief, but I think knowledge is being able to dunk a basketball,” it would then become clear that they were talking about two different things and their discussion would be pointless.

So how close are the laity’s notions of morality to what academic philosophers are studying? Huemer thinks the two are far apart enough that it is not implausible that the two groups are just referring to different concepts. But I think there is good reason to think that they are less far apart than Huemer suggests.

First, even non-philosophers very often appeal to the harm that seedy behaviors cause, either to the agent doing them or to others. Now, it is true that on my view, the harm an act will cause to the agent herself will not ground a moral reason, but that is a substantive view that I am arguing for. Some philosophers reject this view, or do not have a considered view one way or the other. All I am attempting to show right now is that the laity’s notions about morality and those among academic ethicists are not as different as Huemer suggests. At any rate, most will accept that when an act will cause the agent who performs it harm, that tends to be a reason for the agent not to perform the act. If non-philosophers appeal to the harm the behaviors cause in order to justify their claim that the behaviors in question are wrong, this suggests that perhaps they do not really believe that the behaviors are wrong in a basic way, but because of the harms they cause. And this is just the sort of justification professional ethicists might give.

Those in certain religious circles who believe homosexuality (or certain homosexual behavior) is wrong very often claim that it is harmful. Sometimes they appeal to specific harms, and sometimes they simply claim that it is not the life God has planned for an individual.²⁹ Republican Congresswoman Michele Bachman, for instance, said at an education conference in 2004, “Don’t misunderstand. I am not here bashing people who are homosexuals, who are lesbians, who are

²⁹ I hear this line all the time from my in-laws.

bisexual, who are transgender. We need to have profound compassion for people who are dealing with the very real issue of sexual dysfunction in their life and sexual identity disorders.”³⁰ So Bachman seems to think that homosexuality is harmful, and we can assume that the supposed harm explains, at least in part, why people should avoid it.

Now, as Huemer points out, if conservative Christians like Michele Bachman are worried about harm or risk, then they should claim that cigarette smoking or overeating are equally or more immoral than the various and sundry sexual behaviors they are so worried about. Huemer, referring to the risks of sexual promiscuity, writes,

My point here is not simply that those who consider sexual promiscuity immoral are mistaken. My point is that they are employing a concept of morality that appeals to a distinctive kind of reason for action not recognized in contemporary ethics. It is not plausible that they simply have mistaken factual beliefs about such things as the level of risk associated with overeating, or the effectiveness of condoms. Nor is it plausible that they are blatantly inconsistent. Rather, they perceive promiscuity as having a distinctive, negative evaluative property not possessed by such things as taking unnecessary health risks in general, or incorrectly weighing the costs and benefits of one’s choices. (122)

Why people like Bachman believe homosexuality and sexual promiscuity are morally wrong is a complicated question, and to some degree we can only speculate. And the explanation no doubt differs from case to case. I think there are some who really do believe that the wrongness of behaviors like gambling or promiscuity really is explained by the sorts of considerations academic

³⁰ It is interesting that she claims she is not “bashing” homosexuals and that we ought to feel compassion for them. This may suggest that she does not think that homosexuality is not immoral, but simply that there are reasons to avoid it—or least that she might be amenable to arguments for that view.

philosophers appeal to, e.g., the harm the behaviors cause. They might say something like, “Look, I can’t point to every single harm that sexual promiscuity causes. But God who is omniscient and omnibenevolent tells us that we should not be sexually promiscuous. So if we are sexually promiscuous, our lives are bound to go seriously wrong.” I believe there are some who really do reason this way.

But I agree that there are others who do “perceive promiscuity as having a distinctive negative property.” But they might simply be mistaken about this. This alone does not entail that philosophers and conservative Christians mean different things by ‘immoral’ and its cognates. A Kantian may perceive the violation of rights as having a distinctive negative property that a utilitarian does not perceive. The utilitarian thinks rights are a fiction. But this does not mean that their respective theories are theories of two different things.

Imagine the following hypothetical scenario. Julie is raised in a conservative Christian household, and believes everything she is taught in church, e.g., drug use is always morally wrong, sexual promiscuity is morally wrong, divine command theory is true, etc. She then goes off to college and begins taking philosophy classes. One of the first insights she has is that divine command theory is false. She then acquires a keen interest in ethics. Her beliefs about morality become more consistent and refined. She is forced to ask herself, ‘What justification is there for the claim that homosexuality is immoral?’ She comes to doubt that a homosexual lifestyle is harmful, for the person living it or for others. After some reflection, she discovers there is no reason to believe that homosexuality is immoral, and so she gives up that belief. She goes through the same process for gambling, and the use of certain recreational drugs. Very gradually, over the course of several years, she comes to adopt an ethical theory that is widely held among academic ethicists. Let’s say it’s utilitarianism. Her views about morality are now quite different than they were when she was an incoming freshman.

Question: Are Julie's utilitarian views about something other than what her views on morality were about before entering college? It does not seem to me that they are. And it probably would not seem to Julie that they were, either. It is tempting to make the point that because we cannot locate the precise moment when she started referring to a different concept by 'morality,' then it must be the case that she has been thinking of the same concept all along. But I'm hesitant to make such a claim, because it seems that if Julie's ideas about "morality" became at some point sufficiently different from those she previously held, it could (perhaps) be true that she would then be simply thinking about something else. How different the views would have to be to count as being views about different things, I can't say. But intuitively, it seems like she has been talking about the same thing throughout her slow transition from conservative Christian to utilitarian.

Another important consideration to keep in mind is that in fact, there is not wide agreement among non-philosophers that what I have been calling seedy behaviors are morally wrong. There is great public debate about such things as the moral status of drug use, various kinds of sexual activity and (though perhaps to a lesser degree) gambling. And views about these activities are rapidly changing. A recent Gallup poll indicates that 61% of Americans believe having a baby outside of wedlock is morally acceptable, up from 45% in 2001. 68% believe having sex outside of marriage is morally permissible, up from 53% in 2001. And 63% believe homosexual relationships are morally acceptable, up from 40% in 2001. So while it may be true that some do truly believe that these behaviors are morally wrong, this view is far from unanimous. And it does not seem plausible that those in the public arena who hold conservative views about sex and those who take a more permissive view are talking about two different things when they make claims about the moral status of various sexual behaviors, one group about morality, the other about something else.

Finally, consider again that paradigm cases of moral wrongness in academic ethics involve "causing suffering to another being, without that being's consent, and without the expectation of

achieving any compensating value,” while for many non-philosophers, paradigm cases of moral wrongness involve sexual promiscuity. It is true that academic ethicists tend to reject the claim that sexual promiscuity is wrong in itself. But the reverse is not true. Most lay people would think that torturing someone for fun would be morally wrong. Most people do consider Adolph Hitler to be a moral monster. There is wide agreement among philosophers and non-philosophers alike that behaviors like rape, torture, stealing, spousal abuse, bribery, fraud, slavery, racial discrimination, and child molestation are morally wrong. This suggests that while some lay people may “perceive promiscuity as having a distinctly negative property,” having this property is not part of the *meaning* of ‘immorality’ on their view.

3. Conclusion: The Upshot

I have been arguing that a common sense understanding of moral concepts should lead us to a view that moral reasons are exclusively other-regarding. At this point, one might wonder why all this matters. This is especially so because I have claimed that while some actions are not morally obligatory, it is nevertheless the case that an agent still has reasons to perform them, and even ought to perform them. Aren’t we ultimately concerned with what we, to use Owen McLeod’s phrase, *just plain ought* to do (2001)? What do we gain by distinguishing between what we morally ought to do and what we just plain ought to do? I sympathize with this concern, in part because, ultimately, I too am concerned about what we all-things-considered or just plain ought to do.

Yet I still think the topic is important for several reasons. First, getting to the bottom of what is distinctive about moral reasons and moral obligations can help to answer the question about what we just plain ought to do. There has been considerable interest in recent years about whether there is a Just Plain *Ought* or an all-things-considered point of view from which to make practical

judgments, and, relatedly, whether different kinds of practical reasons are comparable.³¹ But in answering these important questions, it will be useful to first get clear on what the practical reasons and *oughts* are in the first place, and what makes them the kind that they are.

Second, the concept of morality, and by extension moral reasons, is apparently of great interest. Philosophers and laypeople alike spend a lot of time thinking about and discussing morality. That by itself seems like a compelling reason to attempt to discover what these concepts are.

Third, it could potentially aid in making certain kinds of normative arguments. Consider again homosexuality. Suppose a person who initially believed that homosexuality was wrong became convinced of the proposition for which I am arguing: That moral reasons are exclusively other-regarding. She would then have to give up the claim that there is something distinctively wrong about homosexual behavior. Suppose we could then convince her also of the descriptive claim that one person's homosexual behavior tends not to negatively affect others in a significant way. She would then have to give up the idea that there is anything at all immoral about homosexual behavior. That would be quite something. Suppose she could be persuaded of a further descriptive claim, that homosexual behavior tends not to be harmful for the person involved in it. She would then have to accept that there is no reason, moral or otherwise, to avoid homosexual behavior

Finally, it can help to settle a longstanding debate in the philosophy of law, which involves the question of whether the state is ever justified in prohibiting immorality as such, that is, prohibiting an act merely on the grounds that it is immoral. Such laws violate Mill's well-known harm principle: "The only purpose for which power can be rightfully exercised over any member of a civilized community, against his will, is to prevent harm to others" (1956, 13). This principle rules out both laws that prohibit immorality as such, as well as paternalistic laws. Contrary to Mill's harm

³¹ See, for example, Copp (1997), McElwee (2007), Mcleod (2001), Portmore (2008), Rachels (2009), and Smith (2009).

principle, Patrick Devlin famously argued that the immorality (or at least perceived immorality) of an action is sufficient to justify prohibition. Today, this view is known as legal moralism.

Most liberals today reject legal moralism. Many take it to be obviously false. For example, many argue that laws that regulate private sexual behavior between consenting adults are unjustified, because the state cannot “legislate morality.” But arguing against legal moralism is more difficult than it may at first appear. As Gerald Dworkin argues, even liberals must accept that there is some important connection between morality and the law (1999). Various attempts have been made to spell out this connection in a way that avoids legal moralism, but the challenge has proven difficult.

If, however, morality only involves our treatment of others, we need not choose between the harm principle and legal moralism since the only actions that are immoral in the first place are those that harm other individuals.³² If the view I have defended is correct, one can embrace Mill’s harm principle *and* legal moralism. It also avoids the implication that, for example, private sexual acts between consenting adults can be legitimately prohibited. Indeed, it seems to me that the underlying motivation for rejecting such laws is not actually that the state cannot be justified in legislating morality, but that such acts are not immoral in the first place. This tidy solution sidesteps the need to identify which immoral acts can be legitimately prohibited. If this view of morality is true, all immoral acts can, at least in principle, be justly prohibited.

³² Admittedly, this requires some qualification. For one thing, I have said that there *may* be ways of wronging people without harming them, but I have not committed either way. If it turns out to be impossible to wrong a person without harming her, then there is clearly no conflict. If, however, there are ways of wronging a person without harming her, say, by failing to honor a contract with her in a way that does not harm her, then it seems like laws against doing that, though they strictly speaking may violate the harm principle as it is stated, still seem to be in the spirit of the harm principle, in part because failing to honor contracts at least *tends* to harm others. Since the law is a blunt instrument, these laws can be justified even if not every instance of failing to honor a contract leads to a harm.

CHAPTER 4

Relational Reasons and Partiality

1. What Are Relational Reasons?

In the last two chapters, I argued that moral reasons are exclusively other-regarding. In this chapter, I will explore the concept of partiality and explain and defend the way in which my view is partial. Before doing this, it will be useful to briefly review the general view I am defending. My general aim in this dissertation is to provide a classification of practical reasons and discuss the role these reasons play in determining what we ought to do. I have argued that there are three basic kinds of reasons:

1. Self-regarding reasons are reasons an agent has to treat herself in certain ways, for her own sake.
2. Other-regarding reasons are reasons an agent has to treat all morally considerable individuals in certain ways, for their own sakes.
3. Relational-reasons are reasons an agent has to treat certain individuals in certain ways for their own sakes, in virtue of a relationship the agent has with those individuals.

Since moral reasons are all and only those reasons an agent has to treat other individuals in certain ways for their own sakes, both other-regarding reasons and relational reasons are moral reasons. Self-regarding reasons, as I have now argued at some length, are non-moral reasons. I remind the reader that although self-regarding reasons are non-moral, they are still morally relevant: though they do not, “morally speaking, count in favor of, or against performing some action” they are nevertheless “relevant to determining an act’s moral status” (Portmore 2008, 371).

What kinds of relationships are morally relevant? There may be several, but here, I am going to focus what I will call *personal relationships*. These kinds of relationships are quite familiar: spouses, children, perhaps other family members, friends, and (at least some) acquaintances. But there may be other kinds of morally relevant relationships, for example, the relationship between a promisee and a promisor, or, similarly, a relationship based on a contract. Likewise, one might have morally relevant reasons based on a role one occupies – perhaps the President of the United States has reasons to prioritize the interests of Americans over the interests of everyone else. It may be that this is a separate kind of reason from one based on a promise, or it may simply be that he has these reasons based on an oath he took. One might also have reasons of gratitude that are in no way connected to any kind of relationship – if someone saves my life while I am unconscious and I never have any conscious interaction with that person, it may be that I owe her a debt of gratitude.

So there may be a variety of different kinds of relationships, some of which are morally relevant, and some of which are not. For my purposes, to say that a relationship is not morally relevant is to say that one should simply reason as a consequentialist when, for example, deciding whether to keep a promise or honor a contract. I remain neutral on whether these other kinds of relationships are morally relevant. I do not, however, remain neutral that the personal relationships of the kind I mentioned above are morally relevant, and the purpose of this chapter is to explore and defend the claim that such personal relationships provide basic practical reasons. One reason I choose to focus on these kinds of relationships is that they seem to me to be especially deep and important to practical reasoning. Another is that, as we will see, they seem especially recalcitrant to consequentialist treatments.

Since this chapter is primarily about the reasons we have to treat those with whom we have a personal relationship in certain ways, in particular, to treat them well, it will be useful to adopt a phrase to refer to such individuals. For the sake of simplicity, I will simply refer to these individuals

as ‘friends.’ This technical sense of the term will include friends, in the everyday sense of the word, but will also include one’s children, spouse, acquaintances, and perhaps other kinds of individuals as well. Later, I will discuss what kinds of relationships are morally relevant.

Another feature of my view that is worth reviewing at this point is how it is that relational reasons yield a partial view. Relational reasons yield a partial view because agents often have more reason to treat their friends well than they do strangers. The most straightforward kinds of examples are cases where the agent stands to benefit other individuals. I should care about the interests of all individuals and I have reasons to make other individuals better off when I can. But because there are relational reasons, I often have more reason to care about the interests of my friends. Practically, this means that it is often permissible to act so as to promote the welfare of my friend, even if I could promote more general welfare by choosing an alternative course of action. In saying this, I do not mean to commit myself to a welfarist position. There may be other ways to treat people well besides by promoting their welfare. For example, there may be general, basic duties of fidelity, that is, duties to tell the truth and keep promises. If there is such a duty, such a duty will be more pressing when it comes to those with whom I have a personal relationship. To see the intuitive appeal of this, we can contrast a case of my, say, lying to my spouse, with a case of lying to a stranger on the street. If someone approaches me and asks me if I have time to hear about the campaign to beautify the neighborhood, I might consider telling a lie by saying, “Sorry, I don’t have any time right now,” when this is not in fact the case. I may have some reason not to lie, but this reason does not seem terribly weighty, and it might be perfectly permissible to tell such a lie in light of other reasons to lie. But there seems to be something importantly different about telling a lie to my spouse. It seems like I have additional reasons not to lie to my spouse. Of course, some of these reasons are welfarist reasons that can be captured by utilitarianism. For example, she may be more likely to find out that I lied, and she may be hurt by my lying, and it may erode her trust in me. But it

seems *prima facie* plausible that I have extra reason not to lie to her even if I have good reason to believe there is very little chance of her finding out, and that reason seems to be *that she is my spouse*. This fact seems to be relevant to my practical reasons in a basic sort of way.

Another way of saying that agents often have reasons to treat friends well that they don't have when it comes to strangers is that relational reasons are often weightier than general other-regarding reasons. What does this mean? In the first chapter, I discussed two potential models, or ways of understanding this claim. I illustrated these two models with an example from Bernard Williams (1981). Suppose an agent sees two people drowning. One is his wife. The other is a stranger. Assume he can save only one. We might say that the agent has equal general other-regarding reason to save each. If those were the only relevant reasons, he would have equal reason to save each. But the agent has an additional reason to save his wife, namely, that she is his wife. So he should save his wife. The problem with this model is that it might be, as Williams put it, "one thought too many." There might be something problematic about the agent reasoning like this: I realize they are both the kinds of individuals that I should save if I can. But since my wife is my wife, I will save her. Instead, we might think that the mere fact that she is his wife is all the reason he needs to save her. So instead, we might prefer a model where the agent's reason to save his wife is simply that she is his wife. The problem with this model is that that isn't really the only reason he has, for he would have reasons to save her even if she were not his wife. I attempted to resolve this dilemma by pointing out that we can only have relationships with the kinds of individuals whom we have reasons to care about in the first place, that is, all morally considerable individuals. So when we mention a relational reason, the general other-regarding reason is already "baked in," so to speak, to the relational reason. For this reason, it makes sense to speak in terms of there being additional reasons agents have with respect to their treatment of those with whom they have a personal

relationship, but to also sometimes say that the relational reasons *just are* the reasons an agent should do what he should do, for to mention a relational reason *assumes* the general other-regarding reasons.

2. Partiality and Impartiality

I have said that my view is partial, and, in a certain, straightforward sense, it is, since, according to it, one is often permitted, and sometimes even required, to show partial treatment to one's friends. However, as Brad Hooker points out, "Many people closely associate morality and impartiality" (2010, 27). Consider a provocative scene from the television show *Game of Thrones*. The young Jon Snow has recently joined 'The Night's Watch, a special order of soldiers charged with protecting the kingdom from outside threats. When they take their vows, they are required to renounce all property and personal commitments, and to promise to have no wives or children. But Jon is in a quandary: his father has been taken prisoner and his brother is leading a rebellion against the current evil rulers, and Jon is considering breaking his vows and abandoning his post at The Night's Watch in order to help them. Jon has the following conversation with a wise elder of The Night's Watch, named Aemon:

Aemon: Tell me. Did you ever wonder why the men of The Night's Watch take no wives and father no children?

Jon: No

Aemon: So they will not love. Love is the death of duty. If the day should ever come when your lord father was forced to choose between honor on the one hand and those he loves on the other, what would he do?

Jon: He would do whatever was right, no matter what.

Aemon: Then Lord Stark is one man in ten thousand. Most of us are not so strong. What is honor compared to a woman's love? And what is duty against the feel of a

newborn son in your arms? Or a brother's smile...We're all human. Oh, we all do our duty when there is no cost to it. Honor comes easy then. Yet, sooner or later, in every man's life, there comes a day when it is *not* easy. A day when he must *choose*.

Is Aemon right when he claims that love is the death of duty? That we must choose between doing what is right and providing special care for our parents, or siblings, or friends? Of course I do not deny that sometimes such dilemmas exist. For example, it would be wrong for me to poison a child to make room for my own child at an elite private school. This would be an impermissible way of behaving partially toward my child.³³

I agree that a moral view ought to be impartial in some sense. Yet there seem to be many cases where partiality is perfectly acceptable, and perhaps even required. For example, it seems wrong to abandon my family to serve the needs of the poor in the third world, even if I could produce more impartial good by doing so.

Thus, on the one hand, we think impartiality is an ideal to strive for. On the other hand, most people see that a certain kind of partiality is often permissible, or even required.

A way of handling this suggests itself: we should distinguish between different senses of the term 'impartial.' Hooker discusses three senses of the term:

1. Impartial application of good (first-order) moral rules
2. Impartial benevolence as *the* direct guide to decisions about what to do
3. Impartial assessment of (first-order) moral rules (26)

I will now discuss these different kinds of impartiality and examine in which of these senses my view is impartial.

³³ Alternatively, consider Susan Wolf's case of a woman whose son has committed a crime and she must decide whether to hide him from the police (253). Though I suspect some may consider hiding him a form of impermissible partiality, Wolf thinks this is a hard case, and I agree. So sometimes it is unclear whether an action is impermissibly partial.

I will begin with (2): “impartial benevolence as *the* direct guide to decisions about what to do.” This is perhaps the most straightforward sense of the term ‘impartial.’ Peter Singer holds such an impartial view. He argues that when a person faces a practical decision, she must “take into account the interests of all those affected by [her] decision.” He continues, “This requires [her] to weigh up all these interests and adopt the course of action most likely to maximize the interests of those affected” (13). He calls this The Principle of Equal Consideration, which

“acts like a pair of scales, weighing interests impartially. True scales favour the side where the interest is stronger or where several interests combine to outweigh a smaller number of similar interests; but they take no account of whose interests they are weighing.” (20-21)

“So,” writes Hooker, if you were impartial in this sense and “you recognized that donating most of your wealth to Oxfam would benefit the starving more than keeping it for yourself would benefit you [or your friends and loved ones], you would donate it” (31). Likewise, if you saw ways to use someone *else’s* resources, perhaps by stealing them, to maximize interests, you would do that as well. Hooker thinks that this implication of this kind of impartiality renders it absurd (31). Interestingly, though Hooker does not endorse this kind of impartiality, he also thinks that the fact that this kind of theft is wrong also explains the wrongness of some kinds of *partial* behavior, in particular, certain forms of nepotism and favoritism, especially when one occupies a certain role. If I’m charged with filling a position in the company that employs me and I decide to give it to my niece simply because she is my niece, I’m commandeering my company’s resources, which are not mine to use in any way I choose. In cases where one occupies a certain role, Hooker thinks that this kind of impartiality is (at least sometimes) an appropriate guide.

Another reason Hooker rejects this kind of impartiality is that it requires “of people a pattern of decisions that effectively deprives them of friendships” (33), since one should only care

about maximizing interests, and not consider whose interests those might be. This is especially relevant since my aim in this chapter is to defend a kind of partiality that permits special treatment to friends and family. I'll discuss similar arguments in a later section.

My view is certainly not impartial in this sense, but that does not necessarily disqualify it as an acceptable view. To insist that a view must be impartial in this sense in order to be acceptable is simply to beg the question against any view other than act-utilitarianism. Whether my view is impartial in sense (1), in the application of good moral rules, depends on what is entailed by impartially applying good moral rules. I take it that impartially applying a good moral rule is doing what the rule says one should do, regardless of who will be harmed or benefitted by applying the rule. Consider the rule, 'the product of a collective enterprise should be divided in proportion to people's contributions to the enterprise' (29). Suppose this is a good rule, and that I have endorsed it. But suppose that I am in charge of distributing the products, and I fail to follow the rule because I realize that if I do, I will get the short end of the stick. In that case, I fail to apply the rule impartially, and clearly, there is something wrong with being partial in this way.

One thing that Hooker notes is that impartially applying good moral rules does not necessarily imply *treating* everyone impartially. According to Bernard Gert, impartial treatment is this:

A is impartial in respect R with regard to group G if and only if A's actions in respect R are not influenced at all by which member(s) of G are benefited or harmed by these actions. (132)

Impartially *applying* a rule need not require impartial treatment because a rule itself may tell us to be influenced by who is harmed and who is benefitted. The above rule involving the collective enterprise is one example. Instead, Hooker says, "Impartial application of a rule consists in being guided solely by the distinctions identified as relevant by the rule" (29). Thus, the following rule can be applied impartially:

When you could devote your own time, attention, or other resources either to benefiting your friends and family or to benefiting people to whom you have no special connection, and when the benefit given would be about the same size, you should choose to benefit your friends or family. (28)

I will call this rule *The Partial Rule*. There is at least nothing incoherent about The Partial Rule, and so there is no reason to think it could not be impartially applied. Of course, we want our rules to be more than coherent. We want them to be true, or good. Determining whether The Partial Rule is a good rule occurs at the level of assessment. Thus, at least in my view, this kind of impartiality turns out to be rather trivial. The much more substantive kind of impartiality, then, is the third kind: Impartial assessment of first-order moral rules.

What is it for a rule to be impartially defensible? At a minimum, it is “to be defensible from an agent-neutral point of view” (35). Evaluating rules from an agent-neutral point of view, according to Hooker, involves “evaluating them apart from any special attachment of yours.” He continues, “In your assessment of rules, you would not give extra weight to benefits that the rules produce for you, your friends, your family, etc.” (35). So, consider again the rule involving the collective enterprise. Failing to *assess* that rule impartially would be to reject that rule purely on the grounds that, if it were followed, I would get the short end of the stick. As Hooker notes, assessing rules in an agent-neutral way has obvious appeal:

In the assessment of rules, agent-neutrality does seem much more appealing than agent-relativity. The agent-relative assessment ‘Everyone’s accepting these rules is good because this maximizes benefits for me’ is utterly unconvincing. So is ‘everyone’s accepting these rules is good because this maximizes benefits for my group’. Agent-neutral assessment effectively eliminates explicit bias towards oneself, one’s group, and indeed anyone with whom one has some special connection. (35)

Hooker later argues that while agent-neutrality is necessary for impartial assessment, it is not sufficient, for if it were, it would entail that elitism of the talented, the view that the talented should get the most benefits, could be impartially endorsed. This is because one could endorse it regardless of whether oneself is talented and would receive a bigger piece of the pie as a result of general acceptance of the rule. But, he argues, elitism of the talented is unacceptable. And Hooker thinks it is better if we can “find a meaning for ‘impartial’ that does make an evaluative stance’s being impartial enough to make it acceptable” (9).

I myself do not share this concern. I prefer to say that a moral stance’s being capable of being impartially endorsed is necessary, but not sufficient, for it to be acceptable. Therefore, I think a rule’s being able to be endorsed from an agent-neutral point of view is necessary and sufficient for it’s being capable of being impartially endorsed. So on my view, because elitism of the talented can be endorsed agent-neutrally, it can be endorsed impartially. Whether it is unacceptable should be argued on different grounds than impartiality.

3. How My View Is Partial

My view can be impartial in the first and third senses that Hooker discusses: Impartial application and impartial assessment of moral rules, respectively, but not in the second - impartial benevolence is the direct guide to decisions about what to do. But as I said above, to simply insist that a view must be impartial in the second sense is simply to beg the question in favor of utilitarianism. These remarks require some qualification, since I am not talking about rules, but reasons. Still, assessing whether one has relational reasons of the sort I’m discussing can be done so impartially. That is to say, one can endorse the view that one has basic reasons prioritize one’s friends in an agent-neutral way. In fact, nearly everyone does endorse it (at least implicitly), *including* those who could be made substantially better off if it was generally rejected. For example, I think a

reasonable person could accept that I chose to invest in a college fund for my child, because he is my child, even though I could have used that money to provide a more fundamental need for her children. And she, in turn, might provide what she can for her children, even though she recognizes that she could produce more impartial good by doing something else. Just as impartially *applying* a rule need not require impartial treatment (in the 2nd sense), since the rule itself may tell us to be influenced by who is harmed and who is benefitted, impartially applying a reason, that is, acting on that reason, need not require impartial treatment, for the reason may be a reason to prioritize the interests of a specific individual.

My view is of course partial in the 2nd of Hooker's senses, since it rejects impartial benevolence as the direct guide about decisions about what to do. On my view, an agent is sometimes permitted to give special treatment to those with whom she has a certain kind of relationship, and that relationship plays a basic role in justifying that treatment. This second part is important because an act-utilitarian could claim that one is often permitted to treat friends differently than one treats strangers, since this can often promote utility. My friend and I may both be better off if we go to lunch together than if we were to each go to lunch with a stranger. But on the most basic level, the reason one should have lunch with her friend is because it promotes utility.

I in fact hold a stronger version of the view than the one I presented above. I believe agents are sometimes *obligated* to provide special treatment to those with whom they have a relationship. This is because I believe that what one ought to do is whatever one has most reason to do (I defend this claim in the following chapter). If one is deciding to perform an action that will benefit a stranger or an action that will benefit a friend, that the friend is a friend will provide additional reason for the agent to benefit the friend, and that means that, all else equal, the person ought to benefit the friend. Of course, the view that one should always do that which one has most reason to do and the view that there are relational reasons can come apart. They do not entail each other and

one can endorse one without endorsing the other. For the most part, in this chapter, I will be defending the more modest version.

At the start of this chapter, I noted there may be a variety of morally relevant kinds of relationships but that I was going to focus on one in particular, namely, *personal relationships*. Again, I am neutral as to whether other kinds of relationships are morally relevant: perhaps a relationship based on a contract is morally relevant, or maybe I should always reason as a consequentialist when deciding whether I should honor a contract.

What then, on my view, constitutes a personal relationship? I cannot give a complete and definite answer. The concept of a personal relationship is by its very nature fuzzy and amorphous. But I take it that we all have an intuitive understanding of the concept and it is often obvious when we have them and when we don't. I have a personal relationship with my spouse, my child, and my best friend. I do not have a personal relationship with Bob, a stranger in Minnesota whom I have never met. Of course, there are many cases around the margins where it is less clear. I may take a taxi ride and strike up a conversation with my driver. Do I now have a personal relationship with my driver? Do I have a personal relationship with a colleague whom I see around the department but hardly ever speak to? Sometimes it will simply be vague. But I think even in these cases, as I suggested earlier, it is plausible that I have some minimal personal relationship with such people, and therefore have minimal relational reasons to treat them well. For example, if I can provide a benefit to my colleague or a stranger, and I have equal impartial reasons to provide the benefit to each, it is plausible that I have slightly more reason to provide the benefit to my colleague, and therefore ought to do so.

As these examples suggest, and is of course obvious, personal relationships come in degrees. On my view, the strength of our relational reasons correspondingly comes in degrees. For example my relational reasons to care about and promote my spouse's well-being, all else equal, will generally

be stronger than my relational reasons with respect to my treatment of my friend, which are in turn stronger than my relational reasons with respect to my treatment of an acquaintance.

One helpful thing I can say about personal relationships is that, for me to be in a personal relationship at all requires some voluntary action on my part. I could avoid having a personal relationship with my taxi driver, for example, by refusing to engage with her at all. This might be rude for impartial reasons, but imagine that I am riding in a limousine and I never choose to lower the partition and chat with my driver. Then the driver is, for all intents and purposes, a complete stranger to me, and I have no relational reasons regarding my treatment of her.

Thus, on my view, relational reasons based on personal relationships are *weakly voluntaristic*. To say that a relational reason is voluntaristic admits of at least three interpretations. One possible view of voluntaristic relational reasons is that one has such reasons only if one chooses to have them. I reject this view. One reason is that, on my view, reasons in general are not voluntaristic in this sense, at least not usually.³⁴ On my view, if we have a personal relationship with someone, we have relational reasons regarding our treatment of that person, regardless of whether we want to have such reasons. I have relational reasons to promote my friend's welfare even if I don't want to have those reasons, and even if I decide I don't care about him anymore. Of course, it may be that he moves away and we lose touch, and our relationship dissolves. In that case, I may no longer have relational reasons regarding my treatment of him. (Or not: maybe I will always have at least a minimal relational reason. I'm neutral on this point.)

A more plausible interpretation of the claim that relational reasons are voluntaristic is that I have them only if I voluntarily enter into a personal relationship. Or, slightly differently, I have them only if I take myself to be in a personal relationship. On this view, though I have relational reasons

³⁴ In the following chapter, I will discuss a view according to which we sometimes have voluntaristic reasons.

whether I want to or not, it is up to me whether I am in the kind of relationship that would generate relational reasons. But I reject this view as well. We can sometimes find ourselves in a personal relationship we did not intend to enter, and in those cases, we will still have relational reasons. Suppose I wait at the same bus stop each morning, and each morning, there is another person waiting at the bus stop. One morning, out of sheer boredom, I begin making idle chitchat with the stranger. Over the course of months, our conversations naturally and gradually become more personal and we get to know each other in a somewhat intimate way. I now have a personal relationship with that person, and have relational reasons regarding my treatment of him, even though I at no time intended to enter such a relationship. If there was an accident at the bus stop, and he and a stranger were pinned under a bus, and I could only save one, I should save him, because I have more reason to save him.

So while I did not intentionally enter into a relationship with him, I voluntarily performed actions that led to that relationship. What kinds of relations does this condition rule out? An important one would be that of mere biological relation. If I discover that there is a child I unknowingly fathered, I do not have reasons of this kind regarding my treatment of her, though I may have a different kind of relational reason. It would also rule out relationships between a hostage and a captor. Though a hostage may get to know her captor quite intimately over the course of her imprisonment, she couldn't avoid those interactions without putting herself in danger. Thus, she has no relational reasons with respect to her treatment of her captor.³⁵

The weak voluntaristic condition seems to capture the right morally relevant personal relationships. It entails we cannot simply be thrust into a personal relationship without our control. But it also entails that we often have morally relevant relationships we did not directly intend to

³⁵ Of course, Stockholm Syndrome may occur, but this is merely a descriptive psychological phenomenon. On my view, no one has any reason to acquire Stockholm Syndrome, and that is as it should be.

have. In fact, many of our relationships are like this, in particular, relationships with our friends – ‘friends’ in the ordinary sense, not in my technical sense. I have never had a friendship ceremony where I made a commitment to be a friend to someone. Friendships often form slowly, and sometimes even unintentionally. I might have a colleague whom I see regularly. We chat at parties and around the office and our relationship slowly grows more intimate. At some point, I may realize that this person has become a friend.

Now, it is plausible that making a commitment increases the moral significance of a relationship. A commitment is itself a voluntary action that can increase the reason-giving force of a relationship. This explains, in part, why marriages are among the most significant relationships we have.³⁶

4. Why Partiality?

So far, I have simply been discussing various forms of partiality and impartiality, and the ways in which my view is partial and impartial. I now turn to defending the way in which my view is partial.

Begin by considering a case, to which I will give the name *Recommend*: The company I work for is hiring, and my friend, Gretchen, has submitted an application. I am considering whether to put in a good word to my boss on Gretchen’s behalf. I know that doing so will give her a significant advantage. Suppose further that I have no evidence that Gretchen is more qualified for the job than any other candidate – I know she would do fine, but I have no idea whether she would be especially good. We can even suppose I know there is a candidate who would even be slightly better. And suppose I’m not considering lying to my boss. I would only tell him that I think Gretchen is capable

³⁶ See Chang (2013) for an illuminating discussion of the normative significance of commitments. I will return to the topic of commitments in the following chapter when I discuss the idea of normative powers.

of doing the job. Still though, I know that even this modest endorsement will give Gretchen an advantage. To put in a good word for her, then, is to treat her partially, since I can assume most of the other candidates do not have an advocate on the inside. Should I recommend my friend?

To say that I should recommend my friend is to say that I should treat her partially. I now turn to defending the kind of partiality I endorse – the kind that says it is permissible to treat your friends partially at the cost of the impersonal good.

A variety of arguments have been made for defending this kind of view. It is common to begin with the premise that it is permissible to have friends (whatever that entails). The question then becomes, is it possible to have friends *and* behave in the ways consequentialism prescribes? Susan Wolf and Michael Stocker both argue that even if this were possible, when a consequentialist acts in a way a friend is typically expected to act, she acts for the wrong reasons. To borrow Stocker's example, if I visit my friend in the hospital, I do not do so ultimately *because* he is my friend, but only because I realize that will maximize good consequences. It is not *his* good I care about, but just the good in general, and it just so happens that visiting my friend in the hospital is a way in which I can promote the good. I will briefly return to these arguments later.

Troy Jollimore takes a different approach. He argues that to be a friend requires that one act in a way that a consequentialist could not act (2000). He too begins with the premise that it is permissible to have friends. He then argues that to be a friend at all requires having certain kinds of feelings or sentiments towards one's friend, which we can describe as simply loving one's friend. The third premise is this: for friendship to exist, one must not only *feel* a certain way about one's friend, a way in which one does not feel about everyone, one must *express* that feeling through action. Interestingly, this premise is inspired by Kagan, who accepts that for friendship to exist, love must be expressed in some way, but denies that preferential treatment is necessary for that expression. Then comes the crux: "there can be no...partiality-free expressions of love" (72).

To see whether this is true, we need an account of what it is for an action to express a feeling. Jollimore expresses this in terms of a counterfactual:

CC: If an action A is an expression of feeling F , then it must be that if F had not obtained, A would not have occurred.

He writes, “For an agent to express her feelings, her feelings must have significant influence over her actions: if she felt differently, she would act differently.” He considers an example:

I express my love for Joan by bringing about certain outcomes which I value: those in which Joan’s desires are fulfilled (preferably by me); those in which Joan and I share enjoyable experiences; those in which I refrain from betraying Joan’s trust, etc. If I felt differently toward Joan, I would not value those outcomes and would not act so as to bring them about. (73-74)

“But,” he goes on, “this presents a problem for the consequentialist agent”:

According to consequentialism, we all ought to agree in ranking outcomes in order of value, and we all ought to bring about the best outcomes we can. Since such an agent’s particular feelings play no significant role in determining which outcomes she is to bring about, she cannot express her feelings through her actions. (74)

One might think that the consequentialist can express love through speech, but as Jollimore notes, if my consequentialist friend tells me she loves me, I would have no reason to believe that she loves me – only that she believes that telling me she loves me promotes the good. Nor does it help to say that a consequentialist should adopt a policy of scrupulous honesty. Not only is such a policy unlikely to actually maximize the good, but the consequentialist who adopts such a policy “has no way to convince other people she has done so” but, as before, her claim that she has done so only gives others evidence that she believes that claiming this is a way of promoting the good. If she

hides the fact that she is a consequentialist, then she cannot explain why, although I am her friend, she is forbidden from showing me any favoritism. Jollimore writes,

Such an agent is caught in a dilemma: she can only justify her conduct to people she cares for by revealing her moral commitments; but once her moral commitments are known, those around her will have no reason to believe her when she tells them she cares for them. (78)

So for Jollimore, for an agent to have a friend requires the agent loving that friend – loving the friend in a way the agent does not love everyone. It also requires expression of that love, which requires action, but this action has to be partial action, since the feelings themselves are feelings of favoring the friend over other individuals.

5. Acting Out

I now want to offer an additional argument in favor of partial treatment towards friends that shares some similarities with Jollimore's argument. Like Jollimore, I assume it is permissible to have friends. Hardly anyone denies that. And, like Jollimore, I think there is something important about the attitudes one has regarding one's friends. For Jollimore, it is simply analytic that friendship requires love. He says, "A friendship cannot exist where feelings of love and affection are not present. If Bill does not like Susan, if he does not feel any warmth toward or attachment to her, then they are not friends" (72).

I agree with Jollimore that friendship requires love in this way. But I now want to make a normative claim about the attitudes one takes towards one's friends: Not only is it analytically true that one loves one's friends, one also has reasons to love ones friends, such that one is *prima facie*

required to love one's friends.³⁷ Now, this might seem odd: if love is required for friendship to exist in the first place – I simply will love my friend if she is in fact my friend– what is the point in saying that I *ought* to love my friend? Is the normative claim not otiose?

It is not otiose. It is one thing to say that if I stop loving my friend, we are no longer friends.³⁸ It is another thing to also say that I have reasons to continue loving my friend, which entails that I have reasons to maintain the friendship. But this is intuitive – most people don't think I should abandon my friendships for no reason. Of course, reasons do not always give rise to obligations. I may sometimes have reasons to stop loving my friend, and thereby end the friendship, for example, my friend might betray me, or I may discover my friend has a deep moral flaw.

To say one ought to love one's friends is to say that one should be partial in one's attitudes towards one's friends. Indeed, if I say I am partial to something, I am expressing that I like it, or love it. To be partial toward someone is to care about that person more than you care about others. As C.S. Lewis says, "To say 'These are my friends' implies 'those are not' (1993, 40-41).

Ought one to love one's friend? Should one be partial in one's attitudes towards one's friends? For most, the answer is obviously 'yes.' But to say something in support of this, think back to Recommend. Suppose I fail to recommend Gretchen. She might ask, "Why did you not recommend me? Didn't you want me to get the job?" It would be one thing to say, "Listen Gretchen, *of course* I wanted you to get the job. But I'm a consequentialist, and I'm therefore not permitted to give you partial treatment." That response would be problematic for the reasons

³⁷ I stated earlier that I was going to use the term 'friend' in a technical sense to refer to all individuals with whom an agent has a personal relationship. I should note here that while Jollimore's analytic claim seems to hold for all friends, in the everyday sense of 'friend', it does not hold for all personal relationships. I may stop loving my child, but that does not by itself entail that that personal relationship has dissolved. And, crucially, even if I stop loving my child, I may still have reasons to love my child, and to behave partially toward my child.

³⁸ I should note that it is probably more accurate to say that we are not friends if I'm not *disposed* to have special feelings for her. If in a moment those feelings vanish, it is not as though our friendship vanished, too.

Jollimore has already noted. But an even *worse* response would be, “Gretchen, frankly, I don’t really care whether you get the job or not. The only thing I care about is that the best person gets the job.” We could expect Gretchen to respond, “Some friend you are.” Interestingly, I think this response helps to reveal the truth of both the analytic claim and the normative claim. In saying this, Gretchen expresses both that I ought to care whether she gets the job, but that if I do not, I am not in fact her friend at all.

I now introduce a general principle. The principle I have in mind is motivated by the simple thought that there seems to be some kind of important connection between the appropriateness of our attitudes and the appropriateness of our actions. The principle is this:

Act Out: If an attitude is justified then any action performed out of that attitude is *prima facie* justified.

If one ought to love one’s friends, and if Act Out is true, then agents are often permitted to show their friends partial treatment. This is because one ought to have partial attitudes towards one’s friends, and acting out of those partial attitudes, which entails giving partial treatment, is often permitted.

In Recommend, it is permissible for me to have Gretchen as a friend. If Gretchen is in fact my friend, it is permissible for me to have partial attitudes toward Gretchen. In this case, it is appropriate that I hope that Gretchen gets the job. What kind of friend would I be if I were indifferent? And if Act Out is true, then it is appropriate, at least *prima facie*, for me to act out of this hope, i.e., to recommend Gretchen for the job.

But is Act Out true? One reason for thinking that it is is that it is very *prima facie* plausible. There is simply something attractive about this kind of connection, or moral harmony, between the moral status of attitudes and that of actions.

If nothing else, we should very much want Act Out to be true. It would be agonizing to hold that morality tells us to have certain attitudes, but also that we must never act out of those attitudes. One might liken morality to a cruel master, holding out a tempting treat to a dog but never letting him have it. Or one might think of morality as being like Lucy from Peanuts, holding out the football for Charlie Brown, only to remove it the moment he tries to kick it. As Stanley Kowalski screams in anguish to Stella, we would scream at morality, “You’re tearing me apart!”

Relatedly, borrowing a term from Bernard Williams, a view that denied the truth of Act Out would be *alienating* (1973), especially in the context of intimate relationships. In a real sense, morality would require that we ignore some of our most deeply held attitudes – attitudes that are intertwined with, and even constitute, our very identity. As Ross noted, I am not simply an agent who stands to promote the good: I am a father, a husband, a friend (1930).

Consider just how strange it would be if morality turned out to be this way. How could it be that certain attitudes are good, desirable, or appropriate, but no action motivated by that attitude is ever good, desirable, or appropriate? It is one thing to say that *some* actions motivated by good attitudes are impermissible. For example, it would be wrong for me to kill all of Gretchen’s competitors in hopes that she gets the job.³⁹ But to say that there are some attitudes that, though themselves good, can *never* be acted on, would represent a schizophrenic view of morality.

But apart from its intuitive appeal, I think there are deeper reasons for thinking Act Out is true. One reason is metaphysical, the other epistemic, though the two are related.

I begin with the metaphysical reason. The kinds of considerations that constitute reasons for having certain attitudes are the same kinds of considerations that constitute reasons for action. That is to say, when I have reasons for having a certain attitude, I will also have reasons for acting in a

³⁹ Interestingly, it may be that in addition to being motivated by the good attitude of concern for my Gretchen’s well-being, this action might also be motivated by some bad attitudes, like callousness or even malice towards Gretchen’s competitors.

way that corresponds with that attitude, because the reasons for having the attitude and the reasons for acting are provided by the very same facts.

Suppose I see a blind woman about to step out into oncoming traffic. I have reasons to hope she does not do this, and I have reasons to prevent her from doing this. And the reason for having the attitude and the reason to perform the action of stopping her from doing this are provided by the same facts, i.e., that she will get hit by a car and be seriously injured if she steps out into oncoming traffic.

Or suppose there are reasons related to desert. I have reasons to hope that a deserving person is rewarded, and I have reasons to do certain things to bring it about that a deserving person gets rewarded. And these reasons are provided by the same facts: that the person deserves the reward.

Or consider justice. That an injustice has occurred is a reason to feel angry. It is also a reason to act out of that anger by trying to correct the injustice that has occurred.⁴⁰

This metaphysical observation leads to an epistemic one: If we grant the metaphysical claim that the reasons that justify certain attitudes and the reasons that justify actions are provided by the same facts, then when we recognize a reason to have a certain attitude, we thereby also recognize (or are at least in a good position to recognize) a reason to act in a certain way. In Recommend, that Gretchen is my friend is a reason to hope she gets the job. If reasons for attitudes and reasons for actions are provided by the same facts, then I have some practical reason to act in a way toward Gretchen that I do not have with respect to her competitors. Provided I have no special reason to help anyone else get the job, and since I *do* have a special reason regarding my treatment of Gretchen, this reason can become a decisive reason in favor of recommending Gretchen.

⁴⁰ See Elizabeth V. Spelman 1989.

There may be a variety of ways of recognizing a reason to have an attitude. One might simply be via intuition. It may simply seem appropriate for me to hope that Gretchen gets the job. Of course, we often do not adopt attitudes by rationally reflecting on what attitudes we have reasons to have. Rather, attitudes often arise spontaneously, without any intention to adopt them. This is particularly true of a certain subset of attitudes, i.e., emotions. There is a robust feminist literature on emotions, and particularly, the relationship between emotion and knowledge. For instance, Alison Jaggar argues that “[E]motions may be helpful and even necessary rather than inimical to the construction of knowledge” (1989, 146). I think this is especially true when it comes to moral knowledge. As Jaggar herself points out, “emotions and values are closely related...If we had no emotional responses to the world, it is inconceivable that we should ever come to value one state of affairs more highly than another” (153).

There are many ways in which emotions and values may be related, but one way they may be related is that having an emotion can be evidence of a reason to have a certain attitude, i.e., that very emotion. This is because emotions are often reasons-responsive, even when they arise spontaneously. Anger, even spontaneous anger, is very often a response to something that is actually worth being angry about, for example, an injustice. As Jaggar writes, “Only when we reflect on our initially puzzling irritability, revulsion, anger, or fear may we bring to consciousness our ‘gut-level’ awareness that we are in a situation of coercion, cruelty, injustice, or danger” (161). I may find that I am angry about a situation, and this can be evidence that I actually ought to be angry,

That an emotion arose spontaneously, as opposed to being the result of rational reflection about what one ought to feel, does not entail that it is not reasons-responsive. Beliefs also often arise spontaneously – I see a tree, and I form the belief that a tree is there. But we should not conclude from the fact that the belief arose spontaneously that I am not responding to my reasons to believe that a tree is there or that this belief has no rational basis.

Because emotions can be reasons-responsive, having an emotion can be evidence that one has a reason to have that emotion. If, as I argued above, reasons for attitudes and reasons for actions are provided by the same facts, then having an emotion can be evidence of a practical reason. So my anger about an injustice can be evidence that I actually ought to do something about the injustice. In cases such as these, emotions can be indispensable in leading me to recognize the moral facts and the duties they entail, for I may never have discovered those facts in the absence of having the emotion in the first place. And because our emotions are often more apparent to us than our practical reasons - they “leap out” at us - our emotions can lead to the recognition of practical reasons we might otherwise have missed. Far from being quirky cases on the margins, it seems to me that this phenomenon is extremely common. If so, then emotions occupy a very important role in the epistemology of practical reasons.

Of course, as Jaggar herself puts eloquently, “Although our emotions are epistemologically indispensable, they are not epistemologically indisputable. Like all our faculties, they may be misleading, and their data, like all data, are always subject to reinterpretation and revision” (163). I can easily grant this. We may sometimes feel angry and then discover that this anger is unjustified, just as we can form beliefs based on our perceptions, but then discover that our perceptive faculties in question are unreliable under the circumstances.

I note that one need not accept the view that emotions can be evidence of a practical reason in order to accept Act Out. One could insist that emotions are not reasons-responsive, and whether an emotion, or any attitude for that matter, is justified must be evaluated on other grounds, and still accept the argument. Act Out only says that *if* an attitude is justified, then some way of acting out of that attitude is also justified.

Are there counter-examples to Act Out? That is, are there cases where a certain attitude would be justified, but no way of acting out of that attitude would be even *prima facie* justified? One

might argue that one is often justified in feeling anger with one's small child – especially if one has been the parent of a small child. But perhaps one is never justified in acting out of that anger. There are (at least) three ways to respond to this concern without rejecting Act Out. First, we could say that being angry with a small child is not justified. Or we could say that it is justified, and that acting out of that anger is also sometimes justified. In order to see whether this second response is plausible, we need to ask what acting out of anger towards one's child entails. Maybe it need not entail hurting the child, or yelling at the child, or in any other way taking one's anger out on that child. It could simply entail leaving the room and punching a pillow. I want to resist this response, for it seems to me that to act on one's anger toward another individual entails some sort of behavior toward that individual. As my earlier comments suggested, anger towards an unjust political regime is not merely evidence that one has a reason to punch one's pillow. Instead, it is evidence that one would be justified in, say, opposing that political regime. Therefore, it seems to me that to act out of anger towards one's child entails some sort of behavior directed at the child – perhaps striking her, or yelling at her. However, while losing one's temper and yelling at one's child is certainly *understandable*, I do not think it is justified in the sense that it is best supported by one's practical reasons.

Thus, I prefer to say that one is in fact not justified in feeling anger at one's small child. Jaggar writes, "Simply describing ourselves as angry...presupposes that we view ourselves as having been wronged" (159). If by 'presupposes' she means 'entails,' then I disagree. I think I can be angry at the bench on which I just stubbed my toe without thinking the bench has wronged me. However, I do think that to say that we are *justifiably* or *rationaly* angry entails that we view ourselves as having been wronged. Assuming that young children are not yet moral agents, I then cannot be justifiably angry at my child. If this sounds odd, I think it is because, as has already been suggested, it is certainly *understandable* for a parent to be angry at one's child, and perhaps even to yell at one's child.

And we sometimes use ‘justifiable’ in this sense, that is, to imply *excusable*. But it is not justifiable in the sense that it is supported by normative reasons. There is a big difference between saying that something is understandable or excusable, as in the case of someone on a diet being tempted by some tasty morsel, and saying that something is justified in the sense that it is well supported by reasons. Of course, there are explanatory reasons why one is angry at one’s child. But it is difficult to see why one would have most normative reason to be angry at one’s small child, or at a bench, for that matter. After all, wouldn’t it be *better* if I could always avoid being angry at my child if my child is not a morally responsible agent?

The principle Act Out, again, is this: If an attitude is justified, then any action performed out of that attitude is *prima facie* justified. I have defended this principle by arguing that when there are reasons for having an attitude, there are corresponding reasons to act out of that attitude, since reasons for having attitudes and reasons for acting are provided by the same facts. And to say that someone has a reason to perform an action just is to say that that action is *prima facie* justified: one is permitted to perform the action, unless there are reasons not to perform it that outweigh the reasons to perform it.

6. Right Reasons

A view that one has a reason to be partial to one’s friends, and that that reason *is* the fact that they are your friends, is appealing because it meets what I will call the ‘right reasons’ requirement. I noted above that Wolf argues that even if a consequentialist could behave toward a friend the way in which we ordinarily expect friends to behave, such a consequentialist agent would be behaving for the wrong reasons. This agent does not favor her friend *because* the friend is her friend, rather, she does it simply because she recognizes she can produce more value by doing so. Consider Stocker’s famous example:

[S]uppose you are in a hospital, recovering from a long illness. You are very bored and restless and at loose ends when Smith comes in once again. You are now convinced more than ever that he is a fine fellow and a real friend-taking so much time to cheer you up, traveling all the way across town, and so on. You are so effusive with your praise and thanks that he protests that he always tries to do what he thinks is his duty, what he thinks will be best. You at first think he is engaging in a polite form of self-deprecation, relieving the moral burden. But the more you two speak, the more clear it becomes that he was telling the literal truth: that it is not essentially because of you that he came to see you, not because you are friends, but because he thought it his duty, perhaps as a fellow Christian or Communist or whatever, or simply because he knows of no one more in need of cheering up and no one easier to cheer up.

Surely there is something lacking here—and lacking in moral merit or value.

(1976, 462)

Interestingly, Stocker uses this example to critique both consequentialism as well as certain kinds of deontology. If Smith's friend asks him why he came to visit, 'because the categorical imperative demands that I come to visit you' is as unacceptable an answer as "because it promotes the good." As Stocker suggests, the right answer, or at least the answer that would satisfy his friend, seems to be simply "because you are my friend." That he is Smith's friend just *is* the reason – both the explanatory reason and the normative reason. And as Stocker notes, we should seek harmony between our motives and reasons, values, and justifications (453).

My view satisfies the "right reasons" requirement in the most direct and straightforward way as any view can. If my friend asks me why I came to visit him in the hospital, my answer, "because you are my friend" is literally and basically true. Of course, if a consequentialist answered this way,

her answer would be true in some sense, but it would be incomplete. For a complete answer, she would have to say something like, “Because you are my friend, I am especially well placed to promote your well-being by coming to visit you.”⁴¹

⁴¹ It may be that sophisticated consequentialism could satisfy this requirement. According to sophisticated consequentialism, one should not always, perhaps not even often, apply the utility calculus to each and every decision (Railton 1984). Instead, one should try to develop character traits and patterns of behavior that will in general promote the good. It may be that if one cultivates love and affection for one’s friends, such that providing care for one’s friends (like visiting them in the hospital) becomes almost second nature, at least in general, conduces to the good. Thus, if you asked your sophisticated consequentialist friend why she came to visit, she could truly, and more completely, answer, “Because you are my friend.” On the one hand, this answer would indeed be more satisfying than “because it promotes the good” because the former has something to do with *you*, in particular. Of course, it is also true that the reason your sophisticated consequentialist friend developed this particular affection for you was because she thought it would promote the good. I’m unsure whether it makes sense to be bothered by this fact, but perhaps sophisticated consequentialism does a fairly good job satisfying this requirement as well.

CHAPTER 5

Self-regarding Reasons

The previous chapter was about relational reasons. I argued that there are reasons grounded in personal relationships that justify giving preferential treatment to those with whom we have a personal relationship. In this chapter, I explore self-regarding reasons. Though these reasons are non-moral, they often justify, and, as we will see, can even require that an agent give preferential treatment to herself.

I will begin my discussion by exploring Roger Crisp's account of, and defense of, the dualism of practical reason, according to which, in addition to a general reason to promote the good, there are agent-relative reasons that sometimes permit agents to pursue their own good at the cost of the overall good. By postulating a dualism of practical reason, Crisp says we can avoid Shelly Kagan's extremism, the view that "any agent at any time is required by morality to perform that act—of those not otherwise forbidden—which can reasonably be expected to lead to the best consequences" (Crisp 1996, 54).

I will then move on to discuss a response to Crisp made by Brian McEwlee (2007). Like Crisp, McElwee also aims to avoid extremism, but, he argues, postulating agent-relative reasons has implausible implications. Instead of postulating a dualism of practical reason, McElwee says we should avoid extremism by claiming that we are not always required to do that which we have most reason to do. As I will try to show, this is not a satisfactory solution, and the dualism of practical reason does not have the implausible implications McElwee claims it has.

1. Agent-Relative Reasons

In his paper “The Dualism of Practical Reason,” Roger Crisp offers a view that “steers between the two extremes” of moralism, the view that all reasons for action are moral reasons, and egoism, “the view that all reasons for action are ultimately self-interested” (53). He begins by considering the views of a modern moralist, Shelly Kagan, who argues for what he calls *extremism*: the view that “any agent at any time is required by morality to perform that act—of those not otherwise forbidden—which can reasonably be expected to lead to the best consequences” (Crisp 54). The moderate view, by contrast, “permits agents to further their own interests at the expense of the overall good.” “These permissions,” writes Crisp, “Kagan describes as *options*” (54). For example, suppose I currently donate 10% of my income to Oxfam. On the moderate view, I may be permitted to spend the remaining 90% on myself and my family “even though this will result in my bringing about less good overall” (54). The defense of the moderate position, according to Kagan, rests on an appeal *cost*. As Crisp summarizes it, “What is said to justify my not sending yet more money to Oxfam is the cost I would incur in so doing. Morality cannot demand any greater sacrifice from me” (55).

Extremism is quite demanding, and, as the name suggests, extreme. To see why, consider an example from McElwee, which involves “someone who gives up most of her spare time and money to helping the distant poor—perhaps donating 50% of her income and spending four nights a week trying to help those less well off than herself—but who nevertheless does less than the most she possibly can” (360).” “Such a person,” he notes, according to extremism, “is acting morally wrong. This flies in the face of some of our strongest moral convictions” (360).

So some, like McElwee, think that the fact that extremism is much more demanding than common sense morality constitutes a reason to reject it.⁴² Despite the demandingness of the view, Kagan thinks it is nevertheless true. To argue for it, Kagan's basic strategy is to show that the moderate is already committed to a reason to promote the good. He does this by considering a number of examples of judgments that can only be made sense of if we have a reason to promote the good. For example, if a building containing a caged bird and a child is on fire, and I can only save one, it seems that it would be wrong to save the bird. Why? Because saving the child is the objectively better outcome. If a reason to promote the good were the only practical reason, then extremism would be true. This is because we would have most reason to promote the most good, and, assuming we are required to do what we have most reason to do, we would then be required to produce the most good that we could, since there would be no practical reasons that recommend doing less than the most good. This places a burden of proof on the anti-extremist, who, in order to resist extremism, "must identify some other reason-giving property the force of which in certain circumstances can at the very least override the force of the property of being maximally good-promoting" (Crisp, 57).⁴³

Identifying this additional reason-giving property is precisely what Crisp sets out to do by proposing what he calls *the dualism of practical reason*, according to which, in addition to a reason to promote the good, there are also reasons grounded in *the personal point of view*. Before explaining this additional source of reasons, I should note that in describing his view as a dualism of practical

⁴² Both Crisp and McElwee seem to be motivated to reject extremism on the basis that it is too demanding. While I too reject extremism, I am not as motivated by the demandingness objection. Rather, it simply appears to me that there are other practical reasons in addition to the reason to promote the good. Even if I am right that there are reasons other than reason to promote the good, morality still might be quite demanding. If in our philosophizing we discover that morality is very demanding, so be it.

⁴³ Crisp later goes on to reject this argument, though he agrees with the conclusion, i.e., that we do have reasons to promote the good. Instead of going into the details of his arguments, I will simply assume that we do have reasons to promote (at least some kinds of) goodness or value.

reason, Crisp does not mean to suggest that these reasons are incomparable. Given what he says later, it is clear that he does not believe, as Sidgwick seemed to, that practical reason is “divided against itself” (1907, 508) and cannot deliver a single, determinate answer about what is to be done, but “makes contradictory demands on action” (Schneewind 374). Crisp thinks the two kinds of reasons can indeed be compared so as to deliver all-things-considered judgments about what is to be done.

As we saw earlier, the defense of the moderate position rests on an appeal to cost. Kagan and Crisp agree that the most promising strategy for justifying the appeal to cost should begin with the recognition of the *personal point of view*. Kagan writes, “Persons have a point of view from which certain objects (the objects of their desires, concerns, interests) take on an importance and are assigned weight disproportionate to the weight an impartial perspective would assign” (1989, 259).

Kagan himself offers and then rejects two strategies of defending options based on the personal point of view. “The first—the *negative argument*—involves the moderate taking a pessimistic, or negative, view of the nature of persons; morality might well require us to promote the good without limit if it could, but it cannot because we cannot be so motivated to do so” (Crisp 58). While Kagan “spends a lot of time arguing that we could be so motivated,” Crisp thinks that “this argument is unnecessary, since we can be said to have certain reasons even if they could not motivate us” (58). I agree with Crisp on this score.

On the second strategy – the positive argument – “it is *desirable* for a moral system to allow agents to pursue their own interest at the expense of the overall good” (Crisp 59). As Kagan characterizes the strategy,

From the [personal point of view] an agent is inclined to give greater weight to his own interests than those interests might merit from the objective point of view. But if the fact that persons are engaged in their subjective standpoints possesses a kind of

moral value in its own right, then there are reasons for the agent to act in keeping with his subjective point of view. That is, there are reasons for the given agent to promote his interests...beyond the level indicated solely by the objective importance of those interests. (333-4, emphasis added)

Crisp agrees with Kagan that this strategy fails since “there does not appear to be any special value that the objective standpoint cannot capture” (59). So Crisp, like Kagan, rejects both strategies for defending options based on the appeal to cost. But he offers a third strategy – the neutral argument – “that allows that all values are assessable from the objective standpoint.” “But,” writes Crisp, “values are to be sharply distinguished, conceptually speaking, from reasons...[E]ven if all values are to be assessed from the objective point of view, it may be that there are certain agent-relative *reasons* which cannot be captured or grounded from that point of view” (60-61).

Why think there are such reasons? Crisp thinks the answer ultimately lies in *the separateness of persons*.⁴⁴ Crisp illustrates his framework with the following thought experiment: “Anna has just been born. There are only three possible lives open to her:

A: A life in which she has flourishing personal relationships, plenty of leisure and a lucrative career, and does nothing for charity.

B: A life in which she has all of these above goods, though not to the same extent because of the fairly large amounts of time and money she donates to Oxfam.

C: A life in which she has none of the above apart from the career, which she pursues solely in order to give all her spare money to Oxfam. (62)

If there were only reasons to promote the impartial good, Anna should live life C. Crisp says, “One might plausibly think, however, that she has stronger reason to live life B. Support for this thought

⁴⁴ This phrase is adapted from Rawls (1971) and later taken up by Scheffler (1982). Crisp discusses subtle differences between Scheffler’s use of this term and his own in a footnote (63).

should come from reference to a non-evaluative reason-giving property which competes with that of being maximally-good-promoting” (62). That property, according to Crisp, “is that these goods are instantiated in *her* life.” He continues,

The fact that the life is hers and that therefore the goods in life *B* will be hers provides her with a counterbalance to the reason to promote the good which justifies and explains her selection of that life over life *C*. The appeal to cost is in fact an appeal to the separateness of persons, which of course importantly shapes one’s personal point of view. I would ask those who are not persuaded by this suggestion to ask themselves the following question: is Anna being *unreasonable* in choosing life *B* over life *C*? (63)

I agree with Crisp’s basic approach. Much has been said in defense of the idea that the *separateness of persons* justifies an agent’s producing less impartial good than she possibly can. In light of the previous chapter, in addition to agent-centered reasons based on the separateness of persons, I think there are other considerations that justify, and even sometimes require, an agent’s producing less good than she possibly can, namely, relational reasons⁴⁵. My goal here is not to review and comment on the literature on the separateness of persons in a detailed way. Rather, I want to summarize Crisp’s approach, register my general agreement, and move on to Brian McElwee’s critique because it raises some very interesting issues, and draws out some important, and perhaps controversial, implications of the view I am defending.

⁴⁵ Crisp apparently agrees that we are permitted to give preferential treatment to those near and dear to us, i.e., our children (61). But he says the reasons that justify such preferential treatment stems from the personal point of view as well. The language he uses suggests that he thinks reasons to give preferential treatment to one’s own children, for example, are a species of agent-centered reasons, while on my view, the reasons one has to treat oneself in a certain way, and the reasons one has to treat those with whom one has a special relationship, are of distinct kinds. Of course, we ourselves often benefit by treating our friends in a partial way, so will often have self-regarding reason to do so.

2. Value, Reasons, and Requirements

McElwee also defends a position that steers between moralism and egoism. He is explicitly motivated by the demandingness objection to extremism. Above I mentioned his example of a person who donates 50% of her income to charity and spends four nights a week helping the badly-off, but could nevertheless do even more to help. He thinks that the fact that extremism entails that this person behaves immorally by not doing even more to help is extremely counter-intuitive, and is therefore a reason to reject extremism. So McElwee, like Crisp, rejects Kagan's extremism, but he critiques Crisp's postulation of a dualism of practical reason, which he attempts to show has very counter-intuitive implications. Further, he argues that

the most plausible way to ground moral options is by appeal to cost to the agent, but not by saying that the cost outweighs what is at stake in either value or in reason-giving force. Rather my suggestion is that the structure of the concept of obligation is such that we are not necessarily obliged to do that which there is most reason to do. (364)

McElwee agrees with Kagan that there is only one kind of reason, and that is a reason to promote the good. But he disagrees with what is at least implied by both Crisp and Kagan, namely, that we are required to do that which we have most reason to do. To return to the Anna case, McElwee thinks that while Anna has most *reason* to choose life *C*, she is not thereby necessarily *required* to choose life *C*. Thus, as I understand it, McElwee's view can be described as a satisficing view about obligation.

We have three concepts, which seem to be importantly related in some way: value, reasons, and requirements. But what, precisely, is that relationship? One simple, straightforward answer is that they are linked in a strong, absolute sort of way: values are taken as primitive, all reasons are grounded solely in value, and all requirements are grounded solely in reasons. According to this

view, we have reasons to promote the good, and only those reasons, and we are required to do that which we have most reason to do. Therefore, we are required to produce the most good that we can. This is Kagan's extremism. To avoid extremism, we can, among other things, either drive a wedge between values and reasons, or between reasons and requirements. Crisp, by postulating non-value-based, agent-centered reasons, drives a wedge between value and reasons, but retains the tight connection between reasons and requirements. McElwee's proposal is to retain the tight connection between values and reasons, and instead drive a wedge between reasons and requirements. The question I now want to consider is: which strategy is better?

We can start by examining a pair of questions:

(1) Action *X* would produce the most value, but do I have most reason to do *X*?

(2) I have most reason to do *X*, but I am required to do *X*?

Are these questions open or closed? If they are both closed then we wind up with extremism, which McElwee wants to avoid. So McElwee must say that at least one of them is open. His arguments can be seen as a way of arguing that (2) is open.⁴⁶ In this section, I will examine and respond to these arguments, and argue that (2) is in fact closed. Then I will offer some reasons for thinking that (1) is open.

Why does McElwee think that (2) is open? That is, why does McElwee think it could be true both that an agent has most reason to do *X*, but it is not the case that the agent is required to do *X*?

McElwee thinks that we cannot be required to perform an action unless it would be *irrational* not to. If it would merely be less-than-fully-rational to refrain, we cannot be said to be required to do it. But why think that? According to McElwee:

⁴⁶ I am unsure what McElwee would say about the openness of (1). He clearly thinks the answer to (1) is 'yes,' but that could be so even if it is open. I will later allude to some of the reasons he thinks the answer to (1) is 'yes' in the next section.

In the case of putative ‘rational requirements’[...]we must[...]ask, ‘Required, on pain of what?’ In order to constitute a demandingness objection, the answer must be something more substantive than ‘Or else one has acted suboptimally’. Rather, it must be something like ‘Or else one is irrational, or unreasonable’. I will argue below that in fact it is not plausible to say that anyone who fails to do what there is most reason to do (or even what they believe, or have reason to believe, there is most reason to do) is either irrational or unreasonable, in the ordinary senses of these words, which each imply some serious personal criticism. (366-7)

The argument that follows begins by noting that we often, perhaps usually, whether on Kagan’s view or Crisp’s view, act in a sub-optimal way, rationally speaking. I agree that we often act in a sub-optimal way. I also agree that it is not plausible to say that one who has acted merely in a sub-optimal, or less-than-fully-rational, way has thereby necessarily acted *irrationally*. As Parfit observes, “The charge ‘irrational’ is at one end of a range of criticisms. It is like the charge ‘wicked.’ We may claim that some act, though not so bad as to be wicked, is still open to moral criticism” (1984, 119). McElwee continues,

If we stray from the ordinary senses of the terms ‘reasonable’ or ‘irrational’, and use either as a philosophical term of art, simply to mean ‘less than optimally rational’, then there can be no demandingness objection to the claim that what there is most reason to do is to live a very self-sacrificial lifestyle.

This is a confusing passage. The thought seems to be that the mere fact that one has done other than what one has most reason to do does not necessarily entail that one has acted *irrationally*, and, further, that we can only be *demand*ed to refrain from acting *irrationally*. To say that we are demanded to refrain from acting in merely a sub-optimal, but not irrational, way somehow stretches

the term ‘demanded’ beyond linguistically intuitive limits. The conclusion then follows: The mere fact that one has most reason to do something does not by itself entail that one is required to do it.

What exactly is supposed to be wrong with saying that we can be required to refrain from acting in a merely sub-optimal way? McElwee later suggests that “demands proper” can only be demands that take the form of “Required, or else one merits some form of serious personal criticism” (367). But why must the criticism be serious? Why can’t it be mild? McElwee appears to be moved by the following consideration:

There is a general worry, raised by Bernard Williams and Susan Wolf, among others, about shoehorning all ethical claims into the mould of obligation or requirement. One concern is that these terms retain their moralistic connotations—the suggestion that failure to comply makes one subject to blame, criticism or punishment—when we speak of requirements in non-moral contexts. Framing all practical questions in terms of demands, permissions and prohibitions can give misleading moral colouring to non-moral claims. (367-8)

If the fact that an agent has most reason to perform an action does not entail that the agent is required to perform that action, what then, if anything, does it entail? According to McElwee,

[v]erdicts from the standpoint of Reason-as-such⁴⁷ are thus better understood not as requirements in any substantive sense, but rather as being more like recommendations. All that is claimed when we say that there is a reason of certain strength to do *X* is that there is this much to be said in favour of doing *X*. Claims about what we have most overall reason to do are simply claims about what there is

⁴⁷ By ‘reasons-as-such,’ McElwee means reasons *simpliciter*, as opposed moral reasons or prudential reasons. It is closely related to McLeod’s postulation of a *just-plain-ought*. He writes, “[McLeod] argues that we can take up a standpoint outside prudence and morality – the standpoint of ‘Reason-as-such’ – and ask what we *just plain ought* to do” (366).

most to be said for doing, not about what we must do, on pain of some sanction. They do not directly imply any demands proper; they do not take the form 'Required, or else one merits some form of serious personal criticism'. It would only be this sort of substantive notion of a requirement or demand which could give rise to a distinctive demandingness objection. (367)

To sum up: To say that one is required to do something implies that one would be irrational for failing to do it. But we are not always irrational for failing to do what we have most reason to do. Sometimes we are only less-than-fully rational. Therefore, we are not always required to do that which we have most reason to do.

I admit that it perhaps sounds strained to say, for example, that one is *required* to choose X over Y even though one has only slightly more reason to choose X over Y , especially in cases where the reasons for choosing X are not moral reasons. But what if we simply ask: *should* one choose X ? If it is stipulated that there is more reason to choose X , the answer seems to be, 'of course.'

Thinking only prudentially, imagine the following scenario: I can Φ today and receive n units of prudential value. Or, I can wait and Φ tomorrow and receive $n+1$ units of prudential value, where each unit of prudential value is only a tiny bit of prudential value. If I were perfectly rational, I would wait and Φ tomorrow. But suppose I don't want to wait and I Φ today. What should we say about my choice? I think we should say I have acted sub-optimally, but only very slightly⁴⁸. *Should* I have waited to Φ until tomorrow? Well, yes. Do I merit serious criticism? No. But I nevertheless should have waited. Why? Because I have more reason to wait. I should have waited, but, colloquially speaking, it isn't a big deal that I did not. But still, I should have.

⁴⁸ It may be that waiting until tomorrow causes me some pain. If that is the case, we can adjust the example so that Φ ing tomorrow results in enough prudential value to make up for the pain. For example, suppose waiting causes me 10 units of prudential disvalue. Then we can adjust the example so that waiting to Φ tomorrow results in $n+10+1$ units of prudential value.

This suggests that at least a variant of (2):

(2*) I have most reason to do X , but *should* I do X ?

is closed.

McEwlee might agree that I should have waited until tomorrow, and that (2*) is closed, but insist that rationality does not *require* that I wait until tomorrow. This entails that there is an important difference between saying “ s should (or ought to) Φ ” and “ s is required to Φ .” Is there?

I am inclined to say there is not an important difference between saying “ s should (or ought to) Φ ” and “ s is required to Φ ,” except perhaps that the latter has a stronger, or more moralistic, *tone*. William and Wolf may be right when they claim that terms such as ‘demands,’ ‘permissions,’ and ‘prohibitions’ may be misleading, but that doesn’t suggest they are wrong or inappropriate. Why not think this consideration simply gives us a reason to be careful not to be misled, to be more precise with our language? We might, for example, say that one is merely *weakly* prohibited, and that one, strictly speaking, should not do that which one is, even weakly, prohibited from doing? If one does, one merits *some* form of criticism, perhaps very mild criticism – in some cases, so mild it may not be worth making.

Or perhaps there is a difference. Perhaps to say that one is required to Φ implies one would be irrational for failing to Φ , while to say that one should Φ implies only that one would be less-than-fully-rational for failing to Φ .

Even so, if McElwee admitted that we are not required to do what we have most reason to do, but we nevertheless *should* do what we have most reason to do, he would be committed to the claim that we *should* produce as much value as we can, since we have most reason to do so. But a view that entails that we should produce as much value as we can is quite demanding, and demandingness is the very thing McElwee was attempting to avoid. To claim that a view is not

demanding because it entails only that one should produce as much good as one can, not that one is required to produce as much good as one can, would hardly be a satisfying response.

We can see the problem more clearly when looking again at McElwee's own example of a person who does quite a lot of good for the badly off – much more than most of us do. McElwee's concern was that the claim that such a person acts wrongly for not doing more flies in the face of some of our firmest moral judgments. But if he claims that a person should produce as much good as she can, he must say that she should be doing more, and this implies that she is not already doing enough.

Even if there is some subtle difference between saying that one is required to do X and saying that one should do X , McElwee's view will still turn out to be quite demanding, since it entails that one should produce as much good as one can. And it seems to me better to focus on what one should do, rather than what one is required to do. The former is more fundamental. The basic question of practical rationality can most succinctly and accurately be characterized as the question of what one, all-things-considered, should, or ought, to do. Therefore, from here on I will be focusing on (2*): I have most reason to do X , but should I do X ?

Above, I began considering an argument that suggests that (2*) is closed. Instead of Φ ing today, I should wait and Φ tomorrow, even if waiting only results in slightly more prudential value. Further support for the idea that (2*) is closed can be seen by considering the commonplace comparison to weight. At risk of stretching the analogy too far, suppose I have two items, and want to determine which is heavier. I place both on the scale and discover item A weighs slightly more than item B . I now have the answer I was looking for. It is true that A does not weigh *much* more, but it weighs more. If the items were gold pieces, and I could only keep one, then I now know which one I ought to keep. We might want to say that, though it would be sub-optimal, it wouldn't be *crazy* to keep B . But when considered in a certain light, it does seem very strange to keep B , for

(assuming all else is equal, e.g., they look roughly the same, I have no special attachment to *B*) I have no reason whatsoever to keep *B* rather than keep *A*. This, I think, does make the choice of *B* appear very odd indeed.

This thought helps us formulate a response to McElwee's suggestion that verdicts from the standpoint of reasons should be understood merely as recommendations. If I am trying to decide between *X* and *Y*, and determine I have more reason to do *X*, what could possibly *recommend* doing *Y* instead? It would have to be some further reason to do *Y*, but we've just said that there is no such reason. If we both know that I have most reason to do *X*, and I choose *Y*, and you ask me why I chose *Y*, I can give no answer to justify or make sense of my decision.

Indeed, the whole practice of moral philosophy or theorizing about what one practically ought to do can be described as offering reasons for doing one action or another. Imagine a contentious philosophy talk where the speaker is presenting arguments for a view that an agent in a given set of circumstances ought to Φ . What kinds of objections will be raised by the audience? By and large, they will consist in offering reasons why it is not the case that such an agent ought to Φ . Suppose the debate has been going on for some time, some offering reasons to Φ , some offering reasons not to Φ . Now suppose someone raises his hand and says, "People, people. Don't you see that this debate is pointless? Why are you all simply assuming that once we know what the agent has most reason to do, we will know what she ought to do?" I would expect the rest of the audience would be rather nonplussed by this line of reasoning. This, I submit, is because determining what one ought to do just *is* determining what one has most reason to do. What else could it be? Several times, McElwee describes postulating a dualism of practical reason (or even pluralism about practical reason for that matter) as *ad hoc* and that the view that all of our reasons are based directly on value is the much more *prima facie* plausible view. But I cannot understand why he thinks that that is the more plausible view, and what I have attempted to do here is show that the view that what we have

most reason to do is what we are required to do is very *prima facie* plausible (indeed, it seems assumed in the very way we debate about practical issues) and that whether we only have reasons to promote the good seems entirely open (though, of course, certainly not obviously false).

I just said that it seems open whether we have reasons only to promote the good. That is, question (1), ‘action *X* would produce the most value, but do I have most reason to do *X*?’ is open. What support can be found for such a claim? One reason to think it is open is that a great many people think the answer is ‘no.’ Indeed, one of the most fundamental and hotly debated questions in moral philosophy can be framed as (1). Consider James and Stuart Rachels’s classic introductory textbook *The Elements of Moral Philosophy*. After explaining consequentialism’s strong intuitive appeal, they more or less take for granted that the consequences of one’s action matter in determining what is right. But they then consider whether anything else *matters* in determining what is right. They say, “If things other than the consequences are important in determining what is right, then [consequentialism] is incorrect” (112). But what is something that is important in determining whether an action is right other than a reason to perform that action? Later, they discuss what they call *backward looking reasons*. For example, the fact that I promised to do something seems to be a reason to do it, and may make it the case that I ought to do it, even if the consequences would be slightly worse than if I broke my promise and did not do it.

There may be not be any reasons other than reasons to promote the good, but whether there are is at least an intelligible, and indeed, worthwhile, question. And identifying such reasons is precisely what Crisp has set out to do.

There is another way we can drive a wedge between reasons and values. In addition to asserting that there are reasons other than value-based reasons, we can assert, as Crisp does, that although I often “have reason to promote the good[...]my reason for promoting the good is *not* that the good is promoted by my so doing” (56). To return to Kagan’s example involving the building on

fire, it may be that I have most reason to save the child, and saving the child would result in the most good, thus, I have a reason to promote the good. However, my reason to promote the good is not necessarily that it would be good-promoting, but that it would, for example, be saving a human life.

Alternatively, it may be that we have reasons to promote only certain kinds of value. To see why, suppose an omniscient and omnibenevolent being, who never lies, appeared to you and told you that glass pig figurines were an objective and important good. Would this give you strong reason to take up glass-blowing? It doesn't seem like it would, even though we know that the proliferation of glass pigs would make the world objectively better. Thus, it may be that we only have reasons to promote certain kinds of value, for example, the value of the welfare of individual beings. As a more plausible example, it may be that beauty is indeed valuable: it may be that a beautiful world is *intrinsically* better than an ugly one, but this may not give anyone a basic reason to produce the impersonal, "free floating" good of beauty if it would not be enjoyed by any sentient being.

This is not to say that we ever have reason to perform an action that will result in *no* value at all. Crisp himself holds that "[a]ny action that results in *nothing* desirable or valuable...can plausibly be said to be unjustified" (56 emphasis added).

There is a final problem with the way McElwee handles value, reasons, and requirements, which is that it can't do all that much to combat extremism. Recall that McElwee is explicitly motivated by the demandingness objection to consequentialism. He wants to maintain that someone who does quite a lot for others - someone who appears to be acting quite praiseworthily from the point of view of common sense - is not acting immorally. To avoid demandingness, he argues that we are not always required to do what we have most reason to do - we are only required to avoid actions that would be decidedly irrational - while maintaining that reasons correspond directly to value. The problem with his approach is that there will be many cases where a person who does

quite a lot for others, and thus is morally admirable from the point of view of common sense, could still do much more. Such a person would then have much more reason to do even more good. My point here is not to criticize McElwee for failing to say just when an action is irrational, as opposed to less than fully rational.⁴⁹ Determining precisely where the line is between less-than-fully-rational and irrational is difficult for anyone to say, and there might often be some gray areas. But a very natural thing to say is that when someone has only slightly more reason to choose *X* over *Y*, choosing *Y* is merely sub-optimal. But when one has *much* more reason to choose *X* over *Y*, choosing *Y* is irrational, and surely, in these cases, one is required to do *X*. Thus, McElwee's view differs from extremism in the following way: when one could produce a small amount of additional impartial value by sacrificing one's own well-being, though one has most reason choose the impartially best action, one isn't required to do so. But when one could produce significantly more impartial value by sacrificing some of one's own well-being, one has much more reason to choose the impartially best action, and is therefore required to choose that action. And it seems that there can still be cases where a person who is morally praiseworthy according to common sense could still produce quite a lot more impartial value by sacrificing still more of one's own well-being. Such a person would then be required to produce more impartial value, but this was exactly the kind of implication of extremism McElwee wanted to avoid.

To sum up so far: It is plausible that there is some connection between values and reasons. The question then becomes: are there any other kinds of reasons. Crisp suggested that there are. He argued that there are special agent-relative reasons that can justify an agent giving priority to her own interests at the cost of the impartial good. Thus, Anna is not required to choose life *C*. She is permitted to choose life *B*. McElwee, who is also interested in avoiding extremism, took a different

⁴⁹ By the same token, neither have I said just when one has most reason to promote the impartial good at the cost one's own well-being, or at the cost of the well-being of one's friends for that matter. I will take this question up later.

tack. He argued that instead of postulating reasons other than reasons to promote the good, we should say that while there are only reasons to promote the good, we are not always required to do that which we have most reason to do. I argued that this is a mistake, since the question, ‘I have most reason to do *X*, but am I required to do *X* (or the slight variant, *should* I do *X*?)’ is closed, while the question, ‘Action *X* would produce the most value, but do I have most reason do *X*?’ is open. Thus, if we wish to avoid extremism either by driving a wedge between values and reasons, or reasons and requirements, we should opt for the latter strategy.

3. The Tyranny of Practical Reasons

We now come to an important, and potentially embarrassing, implication of the view that both Crisp and I hold, though Crisp is perhaps not sufficiently sensitive to it. The view, to be explicit, has two theses: (1) there are reasons other than reasons to promote the (impartial) good and (2) we are required to do the action that there is most reason to do.

To see the potentially troubling implication, return to Crisp’s Anna. Crisp argued, and I agreed, that Anna has agent-relative, or what I’ve been calling self-regarding, reasons to choose life *B* over life *C*, such that she is *permitted* to choose life *B*. But if we are required to do what we have most reason to do, then Anna, in virtue of her self-regarding reasons, is actually *required* to choose life *B* over life *C*.

McElwee raises the concern by pointing out that on this view, Anna may indeed be permitted to live life *B*. Imagining an agent who devotes some portion of her time and resources to promoting the good, but does less than she possibly can because she prioritizes her own prudential value, he writes,

We must ask what we are to say about the agent if she then becomes even more committed to good causes, and does in fact start maximizing the overall good.

Should we say that she has ceased to do what the balance of reasons supports, now that she seems to be ignoring the ‘due’ weight of the agent-relative or egoistic reasons? It sits very uneasily with us to say that now that she is bringing about even more (and significantly more) good overall, she is acting in a way that she has *less* reason to.” (368-9)

And, of course, on the view I just defended, if Anna is acting in a way that she has *less* reason to, she is thereby acting in a way that she *should* not act. That is to say, Anna *should not* choose life C. In other words, not only is Anna permitted to choose life B, she is *required* to live life B.

This may be problematic, not only because it is unintuitive, but it may undercut the very motivation for postulating a dual-source view of reasons in the first place, namely, to provide agent-centered options. But if we are required to do that which we have most reason to do, we may never, or hardly ever, have genuine options. Crisp says, “Ultimately, when the dual-source theorist argues that there are options, she is arguing that there are agent-relative reasons which are stronger than the *pro tanto* reason to promote the good” (72). “But,” replies McEwlee, “this model cannot be right” (371). This is because, in any scenario where I can choose between various actions, X, Y, or Z, it will seldom be the case that the exact same amount can be said in favor of all of them. That is, there will almost always be, perhaps only slightly, more to be said in favor of one than the rest of them. In other words, there will hardly ever be “ties” in terms of the strength of the reasons that support the respective alternatives. The problem then is that if we say that there are agent-relative reasons that are *stronger* than the *pro tanto* reason to promote the good, and we should do what we have most reason to do, then we don’t end up with genuine options. It may be that we could have options in the case of ties, but there are hardly ever ties, so there are hardly ever options.

I agree with McEwlee’s general point here: Since there are seldom ties, one who endorses the view that one is required to do that which one has most reason to do, cannot help himself to

genuine options. To say that we are required to do what we have most reason to do is to espouse a *tyranny of practical reasons*.

Strictly speaking, I accept the tyranny of practical reasons. I accept that on my view, there are never, or hardly ever, genuine options, in the sense that one is required to do what one has most reason to do. And since I think Anna does indeed have additional self-regarding reason to choose life *B*, it may be the case that she should choose life *B*, and should not choose life *C*.

Whether Anna is rationally required to choose life *B* will depend on a variety of details. I admit that, depending on the details, Anna may be rationally, or all-things-considered, required to choose life *B*. And I admit this sounds rather implausible. In the following two sections, I aim to do two things: First, I will argue that even under the tyranny of practical reasons, and even if there are agent-relative reasons, it may not, in fact, be the case that Anna should not choose life *C*. Second, I will explain why, even if the details are such that Anna is required to choose life *B*, this is less implausible than it first appears.

4. Is Anna Required to Choose Life *B*?

I will discuss two reasons why, even under the tyranny of reasons, combined with pluralism about practical reasons, Anna may in fact not be required to choose life *B*.

First, Anna herself may *want* to live life *C*. An agent's desires are relevant to what practical reasons she has, and this is so even if externalism about reasons is true, as I have been assuming it is. This may be because desire-satisfactionism about welfare is true. If I get what I want, I am thereby made better off, and the value of my welfare provides an objective, external reason for me to pursue what I want. If Anna prefers life *C* over *B*, then, in addition to making others better off, she thereby makes herself better off, and now she has even more reason to choose life *C*. If Anna would make

herself and many others better off, and would make no one (significantly) worse off, then she should of course pick life *C*. In this respect, Anna has options of a sort – options based on her desires.

But desires are also relevant to practical reasons if hedonism about welfare is true, since having one's desires satisfied is often pleasurable⁵⁰. To say that Anna might make herself better off in choosing life *C* is not, as one might suspect, an *ad hoc* response, or overstating things. Helping others is often extremely gratifying, and those who devote significant portions of time and energy toward helping others tend to have a much greater sense of well-being than those who do not.

It may be thought implausible that Anna could be well off living life *C*. Whether it is implausible depends to some degree on how we fill in life *C*. Remember, the goal of life *C* is for Anna to produce as much good as she possibly can. As Crisp describes it, in life *C*, Anna has no flourishing personal relationships or leisure, in which she would presumably engage in rewarding projects and hobbies, and pursues her lucrative career solely in order to give all of her spare money to Oxfam. It is an empirical question whether this is the best, or only, way for Anna to produce the most good, and what sort of career would allow her to make the largest contributions to Oxfam. Would Anna produce more good by becoming a corporate lawyer, earning a large salary, and donating it all to Oxfam? Or by becoming a surgeon, moving overseas to a Third World country, and devoting her life to performing life-saving surgeries? And what would be best for Anna? What would be best for Anna depends, at least to a degree, on her own desires, values, and dispositions. It may be that the conveniences and comforts of home are extremely important to Anna, thus, moving overseas would make her very unhappy. On the other hand, it might be that working as a corporate lawyer would be tedious, while saving lives in the Third World would be exciting and allow face-to-face interactions with those whom she is helping. Working as a surgeon might also allow for

⁵⁰ Indeed, it may be that pleasure simply *is* getting what one wants, as opposed to a sensation. If this is right, then, as Chris Heathwood argues, hedonism and desire-satisfactionism amount to the same view.

unusually close and satisfying relationships, both with the people she is helping, and her fellow dogooders, forming close bonds “working in the trenches.” All of this could be extremely gratifying.

Again, all of this depends on certain features about Anna, and this raises another complication with the way Crisp describes the case, namely, that Anna has just been born. But if Anna has just been born, she has no desires or dispositions. But as we have seen, desires are relevant to our own well-being, and therefore, to our practical reasons, all the more so if we have special agent-relative reasons. Presumably the purpose of imagining Anna as having just been born is to consider a “pure” case. The problem is that it might be too pure: we can’t even imagine such a being making a choice like this. Perhaps then it makes more sense to imagine Anna as having just graduated from college and deciding what to do with the rest of her life. She is a mature adult, with her own desires and values. But, to make the case as pure as we can, we can stipulate that she does not have the sort of projects and personal commitments that adults tend to acquire as they live their lives, pursue their careers, etc. It does not seem implausible that Anna *might* desire to live some version of life *C*.

Of course, it is perfectly possible that life *C* is considerably prudentially worse for Anna than many of her alternatives. This may be because Anna does not want to live life *C*. Or it may be because, although she does, there are other factors that are relevant to well-being, like enjoyment, and Anna would enjoy some of her alternatives much more than life *C*. Or it may be that Anna desires *C* to a degree, and would enjoy *C* to a degree, but it would make the satisfaction of other important desires, and certain other significant enjoyments, impossible, such that on balance, though *C* isn’t prudentially terrible, it is prudentially sub-optimal.

Broadly then, there are three possibilities with respect to Anna’s own welfare in life *C*: (1) It could be prudentially very good for her, (2) it could be prudentially very bad for her, or (3) it could be somewhere in between: less than the best, but not terrible. Which of these is true will depend on

certain features about Anna. If (1) turns out to be true, then it is extremely plausible that she has most reason to pick life *C*, and my view entails that she should. If (2) turns out to be true, then the view I am defending may indeed entail that Anna should not pick life *C*, but now this becomes much easier to swallow, since it is not at all implausible to think we are rationally forbidden from choosing lives that would go terribly for us. I will discuss the plausibility of Anna's being required to choose life *B* in the next section. Whether Anna is required to choose *C* if (3) is true depends on many details, in particular, just how much good Anna is producing, and how much of her own well-being she would have to sacrifice. But it may very well be that Anna should indeed choose life *C*: she would have to sacrifice some of her well-being, but she would be improving the lives of others enormously.

Are these the sorts of options that one might want out of a theory of practical rationality? Yes and no. 'No,' because in the case where, in addition to making others better off as a result of choosing life *C*, Anna also makes herself better off, she is rationally required to choose life *C*. Thus, strictly speaking, she has no options. But 'yes' because she is permitted (indeed, required) to do the thing she most wants to do. And one of the primary motivations for options is the conviction that agents are often permitted to do what they want to do.

There is another way in which it might not be the case that Anna is required to choose life *B*, namely, by exercising her *normative powers*. In a manuscript entitled "Do We Have Normative Powers?", Ruth Chang argues that we can sometimes, through an act of will, create reasons. Chang provides a classification of three traditional views about the source of normativity, or reasons. According to source externalism, the source of normativity lies outside of us. I may have a reason to help a stranger in need, regardless of my attitudes. Source internalism, by contrast, locates the source of practical reasons inside of us, i.e., in our desires or motivations. Both externalism and internalism are accounts of what Chang calls our *given* reasons. Externalism is clearly an account of given

reasons: our reasons are given to us, and not provided by us, because they do not depend basically on any fact about *us*. But internalist reasons are also given reasons, since we simply find ourselves with the desires and dispositions we have, and our reasons are just a product of our desires and dispositions. If I desire to Φ , I have a reason to Φ , regardless of my will to have or not have a reason to Φ . According to a third view, source voluntarism, as with internalism, the source of normativity lies “inside of us – but not in our *passive* states such as desiring [...but...] in the *active* state of willing” (4). These will-based reasons “are considerations that are reasons in virtue of some act of will; they are a matter of our creation. They are voluntarist in their normative source. In short, we create will-based reasons and receive given ones” (7).

Chang proposes a fourth view, hybrid voluntarism, according to which we can sometimes create reasons, but, crucially, only when “our given reasons fail fully to determine what we should do” (8).⁵¹ Reasons can fail to fully determine what one should do, or “run out,” in two ways, namely, when “(1) one fails to have more, less, or equal reason to choose one alternative over the other – what we might call a state of ‘equipoise’, or (2) one has most reason to choose one alternative over the other but it is indeterminate how much more – what we might call a state of ‘indeterminate most reason’” (8).

For the purposes of considering Anna, I will be focusing on (1), for two reasons. The first reason is that it is plausible that, in deciding between *B* and *C*, given her agent-relative reasons, Anna’s reasons are in equipoise. The second reason is that when (2) is the case, though one can create a voluntarist reason, that voluntarist reason “cannot change the valence established by [one’s] given reasons” (47).⁵²

⁵¹ Strictly speaking, hybrid voluntarism is open with respect to whether our given reasons are externalist or internalist, but Chang prefers an externalist account of given reasons (11) and so do I.

⁵² Chang offers several arguments for the claim that voluntarist reasons in these contexts cannot change the valence established by given reasons. Among them are that it would make the view

As Chang explains, “Alternatives are in equipoise when they are incomparable, or ‘on a par’ – that is, comparable, but neither is better than the other and nor are they equally good” (8). The possibility that although neither alternative is better than the other nor are they equally good, *but are nevertheless comparable* is crucial here, for I reject incomparability about practical reason. One might object, however, that two alternatives being neither better nor worse than the other, nor equally good, just *is* for those alternatives to be incomparable. However, in her 2002 paper “The Possibility of Parity,” Chang offers powerful arguments for rejecting this view. She argues for the possibility for a fourth relation, that of being *on a par*, where neither alternative is better or worse than the other, nor equally good, but they are nevertheless comparable. Considering her arguments in detail here, however, would take us too far afield.

Chang illustrates how creating voluntarist reasons when given reasons are in equipoise might work by the following example:

Suppose you have been dating Harry on and off for about a year. It turns out that he needs a kidney, and yours would do nicely. If you don’t give him yours, he can get another off the organ donors list, but he would then face a nontrivial risk of rejection. Now suppose your relationship with Harry is such that your reasons to give him your kidney and your reasons to keep your organs intact ‘run out’ – roughly, they fail fully to determine what you have most reason to do. In such cases, I suggest, you can, through an act of will, take a consideration to be a reason and thereby make it one. In particular, by committing to Harry, you take his need for a kidney to be a reason for you to give him yours and can thereby make it a reason. By

susceptible to the fatal flaws of traditional voluntarism, namely, what she calls the Mafioso problem and the regress problem (5-6). See also 46-49.

creating a new reason to give him your kidney you didn't have before, you now may have most reason to give him your kidney. (2)

So on this account, even though agents are required to do that which they have most reason to do, Anna may still have the option to choose life *B* or life *C* by creating a voluntarist reason to choose *B* or *C*, provided her reasons are in equipoise. Are they? They might very well be. Again, it depends on the details. They might be in equipoise because while there is much to be said in favor of choosing life *C*, in virtue of the impartial good she could produce, she also has the agent-relative reasons Crisp postulates. Though comparable, it might be that *B* is not rationally preferable to *C*, and *C* is not preferable to *B*, nor are they equally good.⁵³ In an especially telling passage, McElwee writes, “Crisp urges those unconvinced by his view to ask themselves the following question: ‘Is Anna being *unreasonable* in choosing life *B* over life *C*?’ Perhaps not. But would Anna be unreasonable if, moved by the horrific plight of those in the Third World suffering from disease and starvation, she chose life *C* over life *B*? *Surely* not” (370). That Anna would not be unreasonable for choosing either *B* or *C* has intuitive appeal, and McElwee more or less endorses this claim. His explanation is that we are not required to do that which we have most reason to do. But Chang offers another explanation of how this could be so, namely, that because her given reasons are in equipoise, she can *create* a further reason to do either one, thus making it the case that she has most reason, and therefore ought, to do that which she chooses to do.

That Anna may be permitted to choose *B* or *C* helps us see the underlying motivation for hybrid voluntarism in the first place. As Chang writes, many decisions are *hard* – cases where “the normative relations among the reasons at stake are a highly nuanced and circumstance-sensitive

⁵³ It may be thought I am begging the question by assuming that Anna indeed has agent-relative reasons of the kind Crisp postulates. But that would be to misunderstand the dialectic. I am here attempting to respond to the objection that *if* Anna has these agent-relative reasons, she may be required to choose life *B*, and there is, the objection goes, something unattractive about that implication.

matter [...and...] it is very unclear how to go about determining what those relations are, and the resolution of the case is of great importance” (12). These kinds of decisions include decisions about what career to pursue, whom to marry, whether to have children, and how many, etc. Anna’s case regarding what life path to follow is a prime example of a hard case - highly nuanced and circumstance-sensitive. These are high-stakes decisions, but it is unclear how the pure externalist can explain how reasons balance out as they do, aside from saying “that’s just how things are” (14). Chang writes,

If, however, as hybrid voluntarism might suppose, externalism accounts for the source of only some and not all of our reasons – if it accounts for the normative source of only our given reasons – then voluntarist reasons can step in and explain why the reasons balance as they do, all things considered. Indeed, hard cases are plausibly ones in which given reasons are in equipoise. That is what makes them hard. If there is nevertheless some determinate fact about what we have most all-things-considered reason to do, will-based reasons can explain that fact: the will steps in to fill the gap left by our given reasons. So, for example, in choosing between careers or places to live or people with whom to spend your life, sometimes your given reasons will run out – your given reasons will be in the state of equipoise. Nevertheless, it seems, at least sometimes, that you *can* have all-things-considered most reason to choose what you end up choosing. How is this to be explained? Hybrid voluntarism offers one possible explanation: in such cases, you can, by an act of will, create a new voluntarist reason that favors one alternative over the other, and thereby make it the case that you have most all-things-considered reason to pursue the one thing rather than the other. When explaining why you have most reason to be, say, a philosopher rather than a deep sea diver, we can appeal to your volition; by

an act of will, you have created a voluntarist reason that gives you most all-things-considered reason to be a philosopher. Insofar as there are hard cases in which it is nevertheless true that you have most all-things-considered reason to do one thing rather than the other, voluntarist reasons provide an additional resource with which to fill the explanatory gap left by one's externalist given reasons. In this way, hybrid voluntarism provides the resources – arguably just where they are needed – for avoiding the explanatory shortfall of source externalism. (14-15)⁵⁴

In sum, then, the existence of agent-relative reasons, combined with the view that agents are rationally required to do what they have most reason to do, does not entail that Anna is required to choose life *B*, and forbidden from choosing *C*.⁵⁵

5. What if Anna is Required to Choose life *B*?

In the previous section, I argued that the view I am defending does not necessarily entail that Anna is required to choose life *B*. However, it might, depending on various details, entail just that. In this section, I argue that even if she is required to choose life *B*, this is less counter-intuitive than it might appear. Why, on my view, might Anna be required to choose life *B*?

One reason is that her reasons may not in fact be in equipoise, so she is not in a position to make a voluntarist reason. Of course, it could be that her reasons are not in equipoise because she has most reason to live life *C*. This, as we saw, depends on features about Anna, e.g., her desires, and

⁵⁴ Chang provides a much more detailed defense of the existence of voluntarist reasons than what I've summarized here. See 26-44.

⁵⁵ One might wonder, in light of the moral/non-moral distinction I made earlier, whether voluntarist reasons are moral or non-moral. The answer is simple: If the voluntarist reason you create through an act of will is a reason to treat yourself in a certain way for your own sake, you have created a non-moral reason. If it is a reason to treat another individual in a certain way for its own sake, then you have created a moral reason.

also about what *C* actually entails. I said above that if it turns out that *C* is quite prudentially good for Anna, then she may in fact have most reason, and therefore be required, to choose life *C*.

But it could also turn out that Anna's reasons are not in equipoise because her reasons favor life *B*. This could be the case because she would be prudentially very badly off in life *C*. I mentioned above the challenges of considering Crisp's thought experiment involving Anna, namely, that we are not certain what sort of life would in fact promote the most good, and that it is very difficult to even imagine a person who has just been born making a decision like this. I suggested that in some ways, it might be better to imagine that Anna is a recent college graduate. If Anna is a recent college graduate, then she likely already has a variety of important personal relationships, and living life *C* may require her to abandon those relationships, and so she has additional reason to choose life *B*.

Alternatively, it may be that Anna's reasons are in equipoise, but she makes no commitment to devote her life to helping the worst off, and instead creates a voluntarist reason to do something else.

So in at least some versions of the story, Anna is rationally forbidden from choosing life *C*, because she has most reason to choose life *B*. Is this implausible? I don't think so.

To ask whether Anna is forbidden from choosing life *C* is to ask whether it is possible to be too morally good. In her important paper "Moral Saints," Susan Wolf argues persuasively that the answer is yes. For Wolf, the life of a moral saint would be a life "dominated by a commitment to improving the welfare of others or society as a whole" (420). She argues that one should not aim at being a moral saint, not only because of the sacrifices it would require to one's own welfare, but because it would foreclose "the discovery of and development of significant nonmoral interests and skills" (421). It is worth noting that in addition to thinking that it is possible for one to be too morally good, Wolf seems to (at least implicitly) endorse the view I defended in previous chapters,

namely, that morality is about how we treat others. If that view is correct, to ask whether it is possible to be too morally good is to ask whether it is possible to do too much for others.

In one sense, utilitarians can answer *this* question in the affirmative: One does too much for others when the sacrifices to one's own well-being outweigh the benefits one provides to others. But, according to utilitarianism, this would itself be immoral. (I argued earlier that this is implausible.) But presumably Anna's case is not like this, for if it were, even utilitarians would agree that Anna should not choose life *C*.

One way to see why Anna's being rationally forbidden from choosing life *C* is not implausible is to note that this does not entail that she would be irrational for doing so, only that she would be less than fully rational. Her decision may not merit serious criticism. In fact, it may even merit certain kinds of praise. More on this below. Of course, some will balk even at the idea that choosing life *C* is less rational than choosing life *B*.

But another way to see why this is not implausible is to consider what a disinterested third-party would *advise* Anna to do. This may seem like an odd point to make, since one might assume that this consideration would support the opposite conclusion. Mill obviously thought that such a "disinterested and benevolent spectator" would advise Anna to choose the utilitarian option (1863, 19), that is, life *C*. But think about what kind of advice *actual* disinterested third-parties give, for example, people who write advice columns. When reading advice columns, one doesn't get the sense that the advisors are giving distinctly moral advice. But it would be a mistake to think that they are giving the letter-writers strictly prudential advice, that is, advising them to do only what is best for themselves, and ignoring the interests of others. If that were true, then the advice columnists would not in fact be impartial – they would instead be partial to the interests of the letter writer. But advice columnists often advise their letter-writers to sacrifice their own interests for the sake of morality.

Consider the strongly worded reply of Amy Dickenson of “Ask Amy” to one of her letter writers. This letter writer, Sad Sister, describes a yearly outing she and the other female members of her family take each year. They stay in a hotel, go shopping, and eat at restaurants. Sad’s problem is her sister, who is crushed because she is excluded from these trips. Sad cites several reasons for the exclusion: her sister “doesn’t have very much money” for the outings, she is different from the other family members, and the other family members are “not interested in what she has to talk about.” Sad continues,

“She complains about her aches and pains, and claims to have some kind of neurological disease that some of us feel is more psychosomatic than real and which she uses to avoid getting up for church on Sundays. She also complains about her ex-husband who left her for another woman, but everyone knows it takes ‘two to tango’ and she is not without fault. We’re all very active churchgoers, while she only sporadically attends services...Now she barely speaks to me and has told our relatives that I am a horrible person.”

This is Amy’s reply:

Dear Sad: First, let’s establish that I agree with your sister: You are a horrible person.

Obviously, you can do whatever you want and associate with — or exclude — whomever you want, but you don’t get to do this and also blame the person you are excluding for not “fitting in.”

The only way your sister would ever fit in would be for you to make room for her. You are unwilling to do that, and that is your choice. But her being upset is completely justified, and you’ll just have to live with that.

Perhaps this is something you could ponder from your church pew, because despite your regular attendance, you don't seem to have learned much.

Instead of thinking that advice columnists take themselves to be offering moral advice, or strictly prudential advice, I think it makes most sense to assume that advice columnists take themselves to be offering advice simply about what is to be done, all things considered. Often they give the interests of the letter writers additional weight, but not because they think those are the only interests that should matter to the letter writer, but simply because they assume that one's interests *ought* to get extra weight.

I can only say for myself that, if Anna were asking me for advice, I would advise her to choose life *B*, at least if the details shook out in a certain way. Life *B* seems to balance morality and self-interest in the appropriate way. In it, Anna devotes significant time, energy, and resources toward helping the badly off, while giving due weight to her own interests. I would therefore say to Anna, "Don't sacrifice all of your interests, desires, goals, and projects to helping others. That's too much. I admire you for your concern for the well-being of others, but you should also pursue your own interests and projects."

Again though, it depends on the details. If Anna told me that she strongly *desired* to live life *C*, or that she had good reason to think she would be very well off living life *C*, then I would of course encourage her to do that. But if Anna reported that she believed she would be miserable living life *C*, and that she would only do it out of a sense of moral duty, I would advise her against it. But, as I said before, it is not at all implausible that people, at least in general, should not live lives that make them miserable.

One might worry that if Anna all-things-considered should not choose life *C*, then she would merit some kind of criticism if she in fact chose life *C*. But we must be careful. On my view, Anna does merit some criticism, namely, she is failing to do what she has most reason to do, and

therefore, she acts in a way that is less than fully rational. But to say that one merits some kind of criticism sounds very much like saying one merits *moral* criticism. But Anna surely does not merit any moral criticism. And this leads to the third reason why the claim that Anna should not choose life *C* is less implausible than it sounds. To say that Anna should not choose life *C* is not to say that there is anything *morally* wrong with choosing life *C*. On the contrary, on my view, and according to common sense, Anna would be extremely morally admirable for choosing life *C*. It's just that she is *too* morally admirable.

If Anna chose life *C*, but discovered it made her very badly off, I would say to Anna, "Anna, morally speaking, you are doing great! Good for you! You are very morally admirable. Your heart is in the right place. But you should also not neglect your own interests so much." That this seems like the sensible thing to say to Anna is consistent with, and therefore lends further support to, the claim I defended in the previous two chapters, namely, that moral reasons are reasons to treat others in certain ways for their own sakes. Anna is responding to her moral reasons, and is very morally admirable for doing so. But she is also neglecting her self-regarding, non-moral reasons, and she should spend more time attending to those than she currently is. One reason we may be hesitant to criticize Anna is that so few people are too morally good. Most people need to be encouraged to be morally better.

At this point, one may have lingering worries about the idea that one can be too morally good. To help assuage the worry, let me bring up a certain kind of case considered earlier: A case where one provides a benefit to another at the cost of an even greater benefit to oneself. Here, all should agree that such a person is doing what one has less reason to do, and therefore, all things considered, should not make such a sacrifice. And yet, it sounds really odd to say that such an action is morally wrong. It is true that utilitarianism entails that such an action would be morally wrong, but as I argued earlier, this is just a reason to reject utilitarianism. It is much more intuitive to say, on the

one hand, the action is unusually morally admirable, but still, should not be done. In short, it is too morally good. Of course, Anna's case is not like this. She really would be adding impartial value to the world by choosing life *C*. I make this point only to show that it is not all that strange to say that one can act in a way that is too morally good.

So what should we make of the idea that Anna is rationally forbidden from choosing life *C*? First, whether my view entails this depends on how well off Anna herself would be in life *C*, and this in turn depends on a variety of details, in particular, details about Anna herself. If she herself could be sufficiently well off in life *C*, then she ought to choose life *C*, and that seems exactly right. If she would be miserable in life *C*, then she should not choose life *C*, and that seems right. If the details are such that Anna really shouldn't choose life *C*, because it would make her miserable, this, I have argued, is not implausible. First, the claim that Anna should not choose life *C* does not entail that she would be irrational for doing so, only that she would be less than fully rational. Second, I, as a neutral third party, would not advise Anna to choose life *C*. And finally, it does not entail that Anna would be immoral for choosing life *C*. On the contrary, she would be extremely morally admirable.

6. Conclusion

In this chapter, I have argued that an agent should always do that which she has most reason to do. I have also argued that, because of her self-regarding reasons, Anna is permitted, and perhaps even ought to, choose life *B*, despite the fact that she could produce more impartial good by choosing life *C*. However, I have also claimed that Anna's reasons for choosing life *B* are not moral reasons, and that it is not the case that she *morally* ought to choose life *B*, or that life *B* is morally better than life *C*. These claims lead to potentially worrying questions about the authority of morality and its role in practical rationality. I explore these questions in the next and final chapter.

CHAPTER 6

The Authority of Morality

I have made a distinction between moral reasons and reasons that are merely morally relevant. And I have said that moral reasons are all and only those reasons an agent has to treat other individuals in certain ways for their own sakes, while reasons that are merely morally relevant are the reasons an agent has to treat herself in certain ways, for her own sake. Earlier, I said that a catch phrase for my view is ‘morality is about how we treat other people.’

I now withdraw this catch phrase, or at least want to heavily qualify it. For to many, morality seems to be a special normative domain. In particular, morality appears to many to be authoritative. To say that one morally ought to Φ is usually meant to imply that one really ought to Φ - that one just plain ought to Φ . And often, when describing what the very study of ethics or moral philosophy is, we say that it is the study of what we ought to do, or how we ought to live. This suggests that the scope of morality is rather large – it covers all questions about what we ought to do and how we ought to live.

But if there are non-moral reasons that nevertheless have normative force, one might wonder how morality could be authoritative, since presumably, when it delivers its verdicts, it leaves out considerations, i.e., non-moral reasons, that are actually relevant to what an agent ought to do. And it would not have as wide a scope as is sometimes suggested, since it would sometimes be true that an agent ought to Φ , but not true that the agent morally ought to Φ . Morality, then, does not cover all questions about what one ought to do.

I accept that on my view, morality does not have quite as wide a scope that it covers all questions regarding what to do and how to live. I have already said, and argued that, it is sometimes

the case that an agent ought to Φ , but it is not the case that the agent morally ought to Φ . Not every genuine *ought* is a moral *ought*.

But what about its authoritativeness? In this final section, I want to explain how morality is authoritative on my view. My aim here is not to defend morality's authority. That has been done in many other places.⁵⁶ Instead, because many believe that morality must be authoritative in some sense (and I agree), I want to show how my view is compatible with at least some versions of the authoritativeness of morality.

There are a number of ways in which morality might be thought to be authoritative, or a variety of views about morality's authority. One difficulty in discussing the authority of morality is that there is not uniform language employed in the literature. Sometimes, different writers will use the same term to describe two different views, and, conversely, different writers will use different terms to describe the same view.

1. Moral Rationalism and Moral Overridingness

Two common terms that are often used are 'moral rationalism' and 'moral overridingness.' Alfred Archer describes moral rationalism as the view that "we always have most reason to act in line with our moral requirements (2013, 101). And more formally, "If an act, Φ , is morally required then, from an all things considered normative perspective, Φ -ing is what there is most reason to do (105). Mark van Roojen defines it this way: "The requirements of ethics are requirements of practical reasons" (2010, 495). Douglas Portmore defines it as the view that "agents can be morally required to do only what they have decisive reason to do, all things considered" (2011, 120). And Russ Shafer-Landau defines it as the view that "moral obligations...entail practical reasons" (2003, 190).

⁵⁶ See, for instance, Archer 2013, Hare 1981, 50-61, Shafer-Landau 2003, 190-214, Stroud 1998, van Roojen 2010.

Similar views have gone under different names. For example, Sarah Stroud uses the term ‘overridingness’ to refer to the view that “If S is morally required to Φ , then S has most reason to Φ (1998, 171). Samuel Scheffler uses the same term for the view that “It can never be rational to knowingly do what morality forbids (1992, 53.) And Stephen Darwell uses the term ‘supremacy’ to describe the view that “there is always conclusive reason, all-things-considered, for agents to act as they morally ought” (2006, 286).

But some authors use the term ‘overridingness’ differently. In particular, some talk in terms of *moral reasons* being overriding. Alfred Archer, for instance, defines overridingness this way: “From an all things considered normative perspective, the reasons that support or are provided by moral requirements always override all nonmoral reasons (104).⁵⁷ Héctor Wittwer also defines overridingness as the view that moral reasons are overriding (2015).⁵⁸

It is difficult to make sense of the claim that moral reasons override other reasons. A simple way of understanding the claim is that whenever moral reasons support Φ ing, an agent ought all-things-considered to Φ , regardless of the strength of the agent’s non-moral reasons to Ψ . This view could perhaps be plausible depending on one’s view of moral reasons. For example, suppose one believed moral reasons were welfarist reasons. This view might then imply that when welfare is at stake, one should always do what will maximize welfare, even if there are genuine reasons that support doing some alternative. This view is of course incompatible with my view, but it isn’t preposterous.

Whatever the merits of such a view, this account of overridingness is clearly implausible when combined with the view of moral reasons I have defended. It would entail that if I can

⁵⁷ Archer actually calls this view ‘Weak Overridingness,’ which he distinguishes from Strong Overridingness, which is that view that “from an all things considered normative perspective, moral reasons always override all nonmoral reasons” (104).

⁵⁸ Though elsewhere, he describes it as the view that “there can never be any sufficient reason for not doing what one knows to be morally prescribed (236).

sacrifice the money I was planning to spend on a vacation to prevent a stranger from getting a paper cut, I ought to do that.⁵⁹

Alternatively, Archer suggests that another way of understanding the view is that “the reasons that support moral obligations...override all nonmoral reasons” (104). This is a hard sentence to parse. Is the idea that moral reasons override non-moral reasons only when the moral reasons support a moral requirement? And is this because once the strength of moral reasons reaches a certain threshold, such that one morally ought to perform the act that is morally required, those moral reasons override the non-moral reasons? What then does ‘override’ mean? Stephen Darwall suggests that it means they “defeat (that is, reduce or undermine the force of) other reasons” (2006, 287). But why think moral reasons would be like that? If I have moral reasons to Φ , and am morally required to Φ , why think those reasons to Φ reduce the force of my reasons to Ψ ?⁶⁰

Alternatively, it might mean that reasons that support moral requirements override other reasons in the sense that they outweigh them. Why think that? I cannot see any reason to accept this view unless one thinks one is only morally required to do that which one has most reason to do, which is just the view that I discussed previously, which sometimes goes by the name of moral rationalism, and sometimes by the name of overridingness.

Because of these complications with the term ‘overridingness,’ I will use the term ‘moral rationalism’ to refer to the view that “agents can be morally required to do only what they have

⁵⁹ This example is borrowed from Sarah Stroud, which is itself a slightly modified example from Philippa Foot (180).

⁶⁰ Darwall suggests one way this might go. He writes, “Consider, for example, the pleasure a torturer might take in seeing her victim squirm. In such a case, the wrongness of torture or of taking pleasure in others’ pain seems not only to outweigh any reason provided by the pleasure; it seems to “silence” it” (287). I suppose that might be plausible (though personally, I don’t find it implausible that the pleasure of the torturer still provides him with a reason to torture – he just has much more reason not to). But of course, most cases are not like this. I might have most moral reason to donate \$100 to Oxfam, and I might even be required to do so, but I can’t see why that reduces the reasons I have to keep the money for myself.

decisive reason to do, all things considered” (Portmore, 120), which is logically equivalent to the view Stroud calls ‘overridingness,’ the view that “if S is morally required to Φ , then S has most reason to Φ ” (171) Archer calls it the constraint thesis: “Necessarily if an act, Φ , is not what there is most reason to do then Φ ing is not morally required” (106). Putting it this way suggests an order of explanation: it’s not as though an act’s being morally required makes it the case that one has most reason to do it, and therefore ought to do it. Rather, it is the fact that there is most reason to do it that explains, at least in part, why it is morally required. I think this is the correct order of explanation, but it’s not necessary for the arguments I make here.

Moral rationalism, as I am using the term, implies that the requirements of morality are *real*, that if one is morally required to Φ one is *really*, from the standpoint of Reasons-as-such, required to Φ . I endorse this view. Why think it is true? Archer discusses three intuitions that support this view. The first is that moral requirements place rational constraints on our actions (107). He explains,

If someone has a moral obligation to act in a particular way then we do not think that she is free to choose how to act. Rather, we think that she has most reason to do what is morally required... As many have pointed out, this point seems to be presupposed by our practice of blaming those who freely and knowledgeably act wrongly. To be blameworthy, though, it seems reasonable to think that we must judge that the agent did not have sufficient reason to act as she did. (108)

The second is that “moral requirements provide a rational justification for action” (108). He explains,

If we accept that an act is morally required then there does not seem to be any need to give further rational justification for performing that act...Rationalism provides the perfect explanation for this intuition. The reason why moral requirements serve as

rational justifiers is that moral requirements are always in line with what we have most reason to do. (108)

Finally,

...demonstrating that an act was in line with the balance of reasons serves as a moral justification for action. Showing that an act was in line with what an agent had most reason to do seems sufficient to show that the act was not morally wrong. It would be odd for someone to claim that an act was in line with what she had most reason to do but also morally impermissible. Again, Rationalism is able to provide the perfect explanation for this thought; the reason that this is sufficient moral justification is that moral requirements are always in line with what there is most reason to do. As a result, showing that an act is not what an agent has most reason to do is sufficient to show that it is not morally required. (108-109)⁶¹

2. Moral Rationalism and the JPO

Another way of expressing moral rationalism is in terms of what Owen McLeod calls the Just Plain *Ought*, which he introduces by way of the following example:

Suppose, then, that you find yourself in the following rather difficult situation: what you morally ought to do conflicts with what you ought to do from the perspective of prudence, which conflicts with what you ought to do from the legal standpoint, which conflicts with what you ought to do from the point of view of etiquette, which conflicts with what you ought to do from an aesthetic point of view. In such a situation, there are at least five things that in some sense you should do, but you can

⁶¹ For similar arguments, see Stroud (1998).

do only one of them. Faced with so many conflicting *oughts* or *shoulds* you might ask – or rather exclaim – “What ought I to do?”

This question might appear puzzling, since one might think you already know everything you need to know. You know what you morally ought to do, what you prudentially ought to do, etc. But McLeod suggests that perhaps, in fact, “you do *not* already know everything you need to know. In particular, you do not know what you *just plain ought* to do. The JPO, he explains, is

an *ought* that is not identical to any of the relative or qualified *oughts*...This is not to say that a moral duty, prudential duty, aesthetic duty, or whatever, cannot in some, or even all cases, also be a just plain duty. It is to say that the concept of JPO is distinct from any relative *ought* concept. (273)

McLeod later provides a defense of the coherence of the JPO, which I won’t get into, but here he pauses to provide some support for the idea that it is a distinct concept from various other kinds of *oughts*. He says that when you know your qualified or relative *oughts* and then further ask “But which one ought I, really, to do?” that appears to be a reasonable, that is to say, an open, question. He says it would “not be reasonable if the JPO were simply identical to one of the relative *oughts*” (273).

So, another way of expressing moral rationalism is to say that all moral *oughts* are also JPOs. That is to say, if it is not the case that one just plain ought to Φ , then it is also not the case that the agent is morally required to Φ . This is why Portmore’s claim that non-moral reasons can nevertheless be morally relevant is important. If non-moral reasons were not morally relevant, then morality would not take them into account, and thus, knowing what I was morally required to do would not entail that I knew what I just plain ought to do. But if all moral requirements are also JPOs, then moral requirements *must* take non-moral reasons into account. If the JPO takes into

account all genuine reasons, and all moral requirements are also JPOs, then moral requirements must also take non-moral reasons into account.

This helps to explain something Archer noted earlier: “[d]emonstrating that an act was in line with the balance of reasons serves as a moral justification for action.” If moral requirements are always in line with what an agent has most reason to do, that is, if moral requirements are always also JPOs, then “showing that an act is not what an agent has most reason to do is sufficient to show that it is not morally required.”

Recall an example I used in the first chapter: A friend asks me to give him a ride to the airport, but there is a *Twilight Zone* marathon on television that I want to watch. Assume the friend isn’t badly in need of the ride – it would just be cheaper and more convenient if I gave him a ride. Assume also that the marathon is a rare event, and I won’t have a chance to watch it again in the near future. In this case, it seems I am not morally required to give my friend a ride, and this is because I am not just plain required to give him a ride. My reasons for not giving him a ride, though non-moral, are still relevant to determining the act’s (the act of staying home and watching the marathon) moral status, that is, they make it the case that that act is not immoral. In this way, showing that I do not have most reason to give my friend a ride is a way of morally justifying my action, but this could only be so if non-moral reasons were nevertheless morally relevant reasons.

3. What is a Moral Requirement?

If, as I have been claiming, all moral requirements are also JPOs, if we know that an agent just plain ought to Φ , what do we gain by saying that the agent also morally ought to Φ ? Recall that above, McLeod asserted that the JPO is not identical to any of the relative *oughts*, including moral *oughts*. But he also said that a relative *ought* in a specific case, might also be a JPO. That is to say, it might sometimes be the case both that I morally (or prudentially, etc.) ought to Φ , and that I just

plain ought to Φ . If moral rationalism is true, then all moral requirements are JPOs, but not all JPOs are moral requirements.

What then makes an act morally required? First, as we have seen, it must be just plain required, which in turn entails that it must be the act for which there is most reason to do. What then makes it *morally* required, as opposed to prudentially required, or whatever required? I propose the following definition:

A requirement to Φ is a moral requirement just in case Φ ing is just plain required, and the moral reasons that count in favor of Φ ing are stronger than the combined strength of all the reasons that count in favor of not Φ ing.

In other words, you are morally required to do something if you are required to do it, and you have stronger (or weightier) moral reason to do it than any reason not to do it. So, to consider a simple kind of case, if I can provide a great benefit to someone else at very little cost, though the cost gives me some reason not to do so, my moral reasons for doing so outweigh the reasons for not doing so, so I am morally required to so.

Notice that on this definition, I am not always morally required to do whatever I have most moral reason to do, nor am I always morally required to do an act that I ought to do, and have most reason to do, and the bulk of the reasons for doing it are moral reasons. Consider the following example: Suppose my action X would provide 10 benefit units to Mr. X , and also provide me with 5 units. My other alternative in this case, action Y , would provide 13 units to Ms. Y , but provide me with 0 units. It is plausible that, all else equal, I have most reason to do X , and that I all-things-considered should do X . And it is true that my moral reasons contribute more to X 's being just plain required than my non-moral reasons. Nevertheless, intuitively, since it doesn't seem that it would be morally wrong to do X , I am not morally required to do X ; this definition entails that I would not be, because it is not the case that the moral reasons for doing X are stronger than the reasons for

doing *Y*. And though I have more moral reason to do *Y*, it is not the case that I am morally required to do *Y*, because my moral reasons for doing *Y* are not stronger than the combined strength of my reasons for doing *X*. In other words, I am not morally required to do *Y*, because I am not all-things-considered required to do *Y*. I might choose to do *Y* if I were completely selfless, and if I did *Y*, I would not be open to moral blame, even though, all-things-considered, I had less reason to do *Y*. In this case, I am not morally required to do either *X* or *Y*, but plausibly, I ought to do *X*.

All moral *oughts*, then, are also JPOs. Can the same be said for other kinds of *oughts*, or is there something special about moral *oughts*? I think there are two equally acceptable ways of answering this question. One is to say that the moral *ought* is special: a moral *ought* is special in that it will always entail a JPO, but this is not the case for other kinds of *oughts*. One might think this is true because to say that someone acted immorally is often meant to imply that she did something she (just plain) ought not to have done. As Archer notes, if an act is morally wrong, we don't normally think that any additional explanation for why an agent ought not do it is required. This is perhaps not so with other kinds of *oughts*. Suppose that, from the point of view of etiquette, I ought to Φ . We don't ordinarily think that necessarily I thereby just plain ought to Φ , especially if there are compelling normative reasons not to Φ .

However, I think it is plausible that the reason etiquettical *oughts* do not always imply JPOs is that such *oughts* are not genuine, normative *oughts* at all. This is because reasons to comply with etiquette are not normatively basic – they will always be derivative of some other kind of reason, in particular, the reasons I have laid out. Sometimes I have reasons to comply with etiquette because if I don't, I will embarrass myself, and that will be bad for me. In such cases, I have non-moral reasons to comply with the requirements of etiquette. Or it may be that failing to comply with the demands of etiquette will offend someone, or make someone uncomfortable. In those cases, I have moral

reasons to comply. But in the absence of these reasons, I think I have no real, normative reason to comply with etiquette.

But, since I really do sometimes have genuine, normative self-regarding reasons, it is often the case that a prudential “ought” is a JPO. I have already said that all moral *oughts* are also JPOs. Can the same be said for prudential *oughts* or are moral *oughts* special? I said above that I think there are two equally acceptable ways of answering this. The first was to say that moral *oughts* are special in this way: they are special because they *always* entail a JPO, while other *oughts* – even ones that are supported by genuinely normative reasons – do not. This is plausible because philosophers will often say that although I prudentially ought to do something, I shouldn’t necessarily do it, perhaps because it would be immoral. On this account, some prudentials *oughts* are *not* JPOs.

Alternatively, one might say that we should only ever say that someone ought to do something – any kind of *ought* – if they just plain ought to do it. To see why this might be sensible, recall that often, to say that an act is morally required is meant to imply that one just plain ought to do it. I think the same can perhaps be said of prudence: to say that an act is prudent might also be a way of endorsing it, that is, a way of saying someone just plain ought to do it. In ordinary language, we don’t normally call an act prudent if one all-things-considered ought not to do it, even if we recognize that the person doing it would benefit most from it. Suppose one of the shady loan officers whose loans contributed to the collapse of the housing bubble in 2008 explained his scheme of offering sub-prime mortgages to clients who could never hope to pay them off. Upon hearing it, I rub my chin and say, “Ah, yes, *very* prudent.” “So you like it?” he asks. Scandalized, I exclaim, “Like it? Of course not! Whatever gave you that impression? I hate it!” It would make sense for the loan officer to be puzzled by this response.

Conversely, to say that an act is imprudent is a way of saying someone just plain ought not to do it. After all, if I give my coat to a stranger who needs it more than I do, it sounds strange to call

that act imprudent, even if we recognize that I would have been better off keeping it for myself. In ordinary language, to say that an act is imprudent is almost always meant to imply that an agent (just plain) ought not do it.

On this account, what makes an act imprudent parallels what makes an act immoral:

A requirement to Φ is a prudential requirement just in case Φ ing is just plain required, and the prudential reasons that count in favor of Φ ing are stronger than the combined strength of all the reasons that count in favor of not Φ ing.

On this account of a prudential requirement, we would also need to make a claim that parallels a claim Portmore made about nonmoral reasons: nonprudential reasons, though nonprudential, are still prudentially relevant. Though they do not count, prudentially speaking, in favor of performing an act, they are nevertheless relevant in determining the act's prudential status. Though making loans to people who cannot pay them back is very good for me, it is not prudentially required, since it is not just plain required, since it would be extremely harmful to others. Though the fact that it is extremely harmful to others is not a prudential reason, it is, on this account, still relevant in determining the act's prudential status.

In either case though, the JPO is distinct from the moral *ought* (or the prudential *ought*). The JPO is more fundamental than the moral ought in the sense that a moral requirement is defined in terms of the JPO - there cannot be a moral *ought* without a JPO. To say that the JPO is more fundamental than the concept of a moral requirement is not to diminish the importance of a moral requirement. Indeed, as McLeod notes, the importance of a moral requirement itself can be explained only if there is a JPO (274). To ask whether a moral requirement matters, perhaps in the sense that I should always do what I am morally required to do, is simply to ask whether, ultimately, all-things-considered, from the standpoint of Reasons-as-such, I ought to do that which I am morally required to do.

Similarly, the question “Should I be moral?” only makes sense in light of a JPO. As McLeod and others have pointed out, this question cannot mean ‘Morally speaking, should I be moral?’ for the answer to this question is trivially ‘yes.’ For such a question to be interesting, it must be asking whether I should *really*, that is, ultimately, all-things-considered, from the standpoint of Reasons-as-such, be moral.

On my view, then, morality is authoritative in the way that I think is the most pressing: One is always (just plain) required to conform to the demands of morality. It is never (just plain) permissible to do that which morality forbids.

BIBLIOGRAPHY

- Archer, Alfred. 2014. "Moral Rationalism Without Overridingness." *Ratio* 27 (1): 100-114.
- Bachman, Michele. Presentation to EdWatch National Education Conference, November 6, 2004.
- Benioff, David and D.B. Weiss. 2011. "Baelor." *A Game of Thrones*. HBO. Television.
- Brennan, Neal. 2003. *Chappelle's Show* Episode 1, Season 1. Comedy Central. Television.
- Carse, Alisa L. 1995. "Pornography: An Uncivil Liberty." *Hypatia* 10 (1): 155-182.
- Chang, Ruth. 2013. "Commitment, Reasons, and the Will." In *Oxford Studies in Metaethics, Vol. 8*, edited by Russ Shafer-Landau, 74-113. Oxford: Oxford University Press.
- . 2010. "Do We Have Normative Powers?" Unpublished manuscript.
<http://fas-philosophy.rutgers.edu/chang/10normpowersdist11-30-10.pdf>
- . 2002. "The Possibility of Parity." *Ethics* 112 (4) 659-688.
- Cohen, Carl. 1997. "Do Animals Have Rights?" *Ethics & Behavior* (2): 91-102.
- Crisp, Roger. 1996. "The Dualism of Practical Reason." *Proceedings of the Aristotelian Society* 96: 53-73.
- Darwall, Stephen. 2006. "Morality and Practical Reason: A Kantian Approach." In *The Oxford Handbook of Ethical Theory*, edited by David Copp, 282-320. New York: Oxford University Press.
- Devlin, Patrick. 1958. "Morals and the Criminal Law." <http://faculty.berea.edu/butlerj/Devlin.pdf>
- Dickinson, Amy. 2014. "Sister Exclusion Makes One Sis a Horrible Person." *Ask Amy*, NPR and *The Chicago Tribune*. <http://amydickinson.com/post/85115023030/sisterly-exclusion-makes-one-sis-a-horrible-person>
- Dworkin, Gerald. 1999. "Devlin Was Right." *William and Mary Law Review* 40 (3): 926-946.
- Gendler, Tamar Szabó. 2008. "Alief and Belief." *Journal of Philosophy* 105 (10): 634-663.
- Gert, Bernard. 1998. *Morality*. New York: Oxford University Press.
- Glasgow, Joshua. "It's Easy Being Me: Duties to Self." *Pea Soup* (blog), August 26th, 2004,
http://peasoup.typepad.com/peasoup/2004/08/its_easy_being_.html
- Hare, R.M. 1981. *Moral Thinking*. Oxford: Oxford University Press.
- Heathwood, Chris. 2006. "Desire Satisfactionism and Hedonism." *Philosophical Studies* 128 (3): 539-563.

- Hooker, Brad. 2010. "When is Impartiality Morally Appropriate?" In *Partiality and Impartiality: Morality, Special Relationships, and The Wider World*, edited by Brian Feltham and John Cottingham, 26-41. Oxford: Oxford University Press.
- Hill, Judith M. 1987. "Pornography and Degradation." *Hypatia* 2 (2): 39-54.
- Huemer, Michael. 2009. "Values and Morals: Outline of a Skeptical Realism." *Philosophical Issues* 19 (1): 113-130.
- Hurley, Elisa A. and Colleen Macnamara. 2011. "Beyond Belief: Toward a Theory of the Reactive Attitudes." *Philosophical Papers* 39 (3): 373-399.
- Jaggar, Alison M. 1989. "Love and Knowledge: Emotion in Feminist Epistemology." *Inquiry* 32 (2): 151-176.
- Jollimore, Troy. 2000. "Friendship Without Partiality." *Ratio* 13 (1): 69-82.
- Jones, Jeff and Lydia Saad (2015) Gallup poll social series: values and beliefs. *Gallup News Service*.
http://www.gallup.com/file/poll/170798/Moral_Acceptability_140530.pdf
- Joyce, Richard. 2008. "Morality, Schmorality." In *Morality and Self-Interest*, edited by Paul Bloomfield, 51-75. New York: Oxford University Press.
- Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Oxford University Press.
- Lewis, C.S. 1993. "Friendship – The Least Necessary Love." In *Friendship: A Philosophical Reader*, edited by Neera Kapur Badhwar, 39-47. Ithaca: Cornell University Press.
- MacKinnon, Catherine A. 1989. "Sexuality, Pornography, and Method: 'Pleasure Under Patriarchy.'" *Ethics* 99 (2): 314-346.
- McElwee, Brian. 2007. "Consequentialism, Demandingness and the Monism of Practical Reason." *Proceedings of the Aristotelian Society* 107 (Part 3): 359-374.
- McLeod, Owen. 2001. "Just Plain 'Ought.'" *Journal of Ethics* 5 (4): 269-291.
- Mill, John Stuart. (1859). 1956. *On Liberty*. Indianapolis: Bobbs-Merrill. Citations refers to the Bobbs-Merrill edition.
- Mill, John Stuart. (1863) 2001. *Utilitarianism*. Kitchener: Batoche Books Limited. Citation refers to the Batoche edition.
- Norcross, Alastair. 2006. "The Scalar Approach to Utilitarianism." In *The Blackwell Guide to Mill's Utilitarianism*, edited by Henry West, 217-232. Oxford: Wiley-Blackwell.
- Parfit, Derek. 2001. "Rationality and Reasons." In *Exploring Practical Philosophy: From Action to Values*, edited by Dan Egonsson and Ingmar Persson, 17-39. England: Ashgate.

- . 2011. *On What Matters*. New York: Oxford University Press.
- . 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Portmore, Douglas W. 2008. "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11 (4): 369-388.
- . 2011. "Consequentialism and Moral Rationalism." In *Oxford Studies in Normative Ethics Volume 1*, edited by Mark Timmons, 120-142. Oxford: Oxford University Press.
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13 (2) 134-171.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.
- Ross, William David. (1930) 2002. *The Right and the Good*. Oxford: Clarendon Press
- Shafer-Landau, Russ. 2003. *Moral Realism*. Oxford: Clarendon Press.
- Scheffler, Samuel. 1992. *Human Morality*. New York: Oxford University Press.
- . 1982. *The Rejection of Consequentialism*. Oxford: Oxford University Press.
- Schneewind, J.B. 1977. *Sidgwick's Ethics and Victorian Moral Philosophy*. Oxford: Clarendon Press.
- Sidgwick, Henry. 1907. *The Methods of Ethics*, 7th edition. Indianapolis: Hackett.
- Singer, Marcus G. 1963. "On Duties to Oneself." *Ethics* 69 (2): 202-205.
- Singer, Peter. 2011. *Practical Ethics*. New York: Cambridge University Press.
- Skorupski, John. 1999. *Ethical Explorations*. Oxford: Oxford University Press.
- Slote, Michael, 1984. "Morality and Self-Other Asymmetry." *Journal of Philosophy* 81 (4): 179-192.
- Smith, Angela M. 2007. "On Being and Holding Responsible." *Journal of Ethics* 11 (4): 465-484.
- Spelman, Elizabeth V. "Anger and Insubordination." In *Women, Knowledge, and Reality: Explorations in Feminist Philosophy*, edited by Ann Garry and Marilyn Pearsall, 263-274. New York: Routledge.
- Stocker, Michael. 1976. "The Schizophrenia of Modern Ethical Theories." *Journal of Philosophy* 73 (14): 453-466.
- Stroud, Sarah. 1998. "Moral Overridingness and Moral Theory." *Pacific Philosophical Quarterly* 79 (2): 170-189.

- Tirrell, Lynne. 1999. "Pornographic Subordination: How Pornography Silences Women." In *On Feminist Ethics and Politics*, edited by Claudia Card, 226-243. Lawrence: University Press of Kansas.
- van Roojen, Mark. 2010. "Moral Rationalism and Rational Amoralism." *Ethics* 120: 495-525.
- van Zyl, Liezl. 2010. "Motive and Right Action." *Philosophia* 38 (2): 405-415.
- Williams, Bernard. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For Against*, by J.J.C. Smart and Bernard Williams, 75-150. Cambridge: Cambridge University Press.
- . 1981. "Internal and External Reasons." In *Moral Luck: Philosophical Papers*, 1-19. Cambridge: Cambridge University Press.
- . 1981. "Persons, Character, and Morality." In *Moral Luck: Philosophical Papers*, 1-19. Cambridge: Cambridge University Press.
- Wikipedia, s.v. "Dwarf-tossing," last modified March 21st 2016,
http://www.chicagomanualofstyle.org/16/ch14/ch14_sec248.html
- Wittwer, Héctor. 2015. "Is the Overridingness of Moral Reasons a Semantic Fact?" In *Why Be Moral?* edited by Beatrix Himmelman and Robert Loudon, 235-248. De Gruyter.
- Wolf, Susan. 1982. "Moral Saints." *Journal of Philosophy* 79 (8): 419-439.
- . "Morality and Partiality." *Philosophical Perspectives* 6: 243-259.
- . 2012. "One Thought Too many: Love, Morality, and the Ordering of Commitment." In *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams*, edited by Ulrike Heuer and Gerald Lang, 71-92. Oxford: Oxford University Press.