

SEPARATING REASONS

by

David Dexter

Submitted in partial fulfilment of the requirements
for the degree of Master of Arts

at

Dalhousie University
Halifax, Nova Scotia
August 2013

© Copyright by David Dexter, 2013

TABLE OF CONTENTS

ABSTRACT	iii
ACKNOWLEDGEMENTS	iv
CHAPTER 1: INTRODUCTION	1
CHAPTER 2: THE DUALISM OF PRACTICAL REASON	9
CHAPTER 3: THE OVERDEMANDINGNESS OF MORALITY	13
CHAPTER 4: PRACTICAL INCONSISTENCY	19
CHAPTER 5: HURLEY’S INTERPERSONAL STANCE	26
CHAPTER 6: DENYING THE IMPERSONAL	34
CHAPTER 7: COMMUNITARIAN MORALISM	40
CHAPTER 8: COMMUNITARIAN MORALISM AND MORAL MAXIMIZATION	60
CHAPTER 9: IT TAKES A VALUE TO MAKE A VALUE	70
CHAPTER 10: CRITICIZING AGENTS’ ACTIONS	75
CHAPTER 11: CONCLUSION	82
BIBLIOGRAPHY	87

ABSTRACT

When facing a dilemma about what to do, rational agents will often encounter a conflict between what they ought to do, morally speaking, and what they most want to do.

Traditionally we think that when there is a moral imperative for an agent to do something, even if she does not want to do it, she nevertheless ought to do it. But this approach inevitably fails to be able to explain why agents often choose to do what they most want, in many cases flouting such moral imperatives. The purpose of this thesis is to offer a plausible alternative to this way of understanding these deliberative dilemmas. I argue that *communitarian moralism*, the account according to which genuine moral imperatives are only imperatives on communities, rather than agents, and according to which agents' moral conduct is necessarily bound up with her particular preferences, projects and commitments, is the most plausible way to understand dilemmas in which agents must choose between doing moral and self-interested actions.

ACKNOWLEDGEMENTS

There are many people that have helped me to get this thesis to its final stages. First I would like to thank my supervisor Prof. Duncan MacIntosh as well as my second reader Prof. Greg Scherkoske and my third reader Prof. Kirstin Borgerson whose comments and support have been invaluable to me. I would also like to show my appreciation for Dalhousie University's philosophy department as a whole; it is hard to find a better atmosphere to immerse oneself in philosophical thought. I want to thank my family and friends both inside and outside philosophy for their continued support of my academic pursuits. Finally, I would like to thank Katie Stockdale for all the personal and philosophical support that she has given me as this project took shape. I am deeply indebted to her for the many brainstorming sessions when I was philosophically stuck and for her motivating words just when I needed to hear them.

CHAPTER 1: INTRODUCTION

Practical dilemmas can arise for an agent when what it is rational to do conflicts with what it is moral to do. The standard account of such dilemmas holds that when an agent chooses to act she carries out a practical deliberation, weighing her reasons for acting rationally against her reasons for acting morally; the outcome of this deliberation is the resolution of her dilemma and she then performs whichever action wins out. That is, the outcome of her deliberation both decides and motivates her to do the action it prescribes. Throughout this project I call rational actions those actions that conduce to the realization of the acting agent's particular projects and that uphold her particular commitments.¹ Rational actions in the foregoing sense are therefore the agent's self-interested actions because, by definition, she cares about the projects and commitments that will be advanced by her acting rationally. I call moral actions those actions that, when performed, will help to advance some specified moral project and moral projects are just those projects that are required by the demands of morality.² So, what the particular moral project is will be determined by what sort of ethics and metaethics one endorses, or perhaps which ethical and metaethical theories turn out to be right.

When an agent does the moral thing, we say that her reasons to act morally outweighed her reasons to act rationally.³ Conversely, if she does the rational thing, her

¹ I do not mean one's 'particular' commitments in contrast to one's general commitments, rather I mean 'particular commitments' in the sense that we might also call them 'discrete commitments'. One's particular commitments are just commitments to the specific things that one cares about, whatever their content happens to be.

² There will be overlap between moral actions and rational actions: some actions are both moral and rational. But moral actions are those actions that must be done whether or not one cares about doing the moral thing, whereas actions that are rational are so just because one cares either about the outcome of performing them or about performing the action itself.

³ My project assumes that one's self-determined actions will be the ones that issue directly from one's having a balance of reasons to so act. Put differently, I assume that when one genuinely chooses to act morally, it is because one has a balance of reasons in favour of doing the moral thing. I also assume here

reasons to act rationally outweighed her reasons to act morally.⁴ While it is possible for the agent to make a mistake in grasping the particular moral and rational reasons for and against an action, once she reaches the end of her deliberation she cannot mistakenly act against the reason that wins out.⁵ On this account, moral and rational reasons are commensurable; they are each reasons for an agent to act that can be weighed against the other. This is interesting for the moral theorist because it can give an explanation for how one answers the question “Why should I be moral?” One should be moral just in case one has better reason to act morally than one has to act rationally.

This line of thought can be found in its early stages in Henry Sidgwick’s *The Methods of Ethics*. Sidgwick famously showed what has since been called the dualism of practical reason. Practical reason, he thought, was divided between those reasons for action that stem from the utilitarian maxim of rational beneficence and those that are derived from the egoistic maxim of prudence.⁶ Because of this dualism, one can contrast one’s reasons for acting as a good utilitarian with one’s reasons for acting as a good egoist. Although Sidgwick thought that a conflict between the two sides would result in

that one’s reasons to act morally can outweigh one’s reasons to act rationally whether one cares about the moral project or not. The reasons to act morally carry their weight, at least on this standard account, just because they are bound up with the *moral* project. We will return to why affording moral reasons special significance in virtue of the fact that they are moral is implausible in chapters six and seven.

⁴ This does imply that there are sometimes reasons to act irrationally. This amounts to little more than saying that sometimes agents have reasons to do things that they do not care about doing. But there is, I think, something unsatisfying about this claim. One of the consequences of the communitarian framework that I put forth in chapter seven is that agents no longer have reasons to act irrationally.

⁵ For example, I might deceive myself into thinking that the satisfaction of a self-destructive desire is better for me than the satisfaction of my prudential desire to not succumb to self-destructive desires. When I try to decide whether I should satisfy the self-destructive desire I will make a deliberative mistake because I hold false beliefs about what I most care about. But once my deliberation is complete, I cannot mistakenly intentionally act against the outcome of my deliberation. If I do, it is only because I have realized that I care more about satisfying my prudential desire, that I had some other false belief in my deliberation or that I made some other deliberative mistake. But by recognizing one of these features, I have not just mistakenly and intentionally acted against the outcome of my initial deliberation. Instead, I have entered a new deliberation, one that involves reviewing my former deliberation to see if I had made a mistake.

⁶ Henry Sidgwick, *The Methods of Ethics*, ed. Jonathan Bennett (2011) accessed April 25, 2013, <http://www.earlymoderntexts.com/pdf/sidgmeth.pdf>, 243.

practical paralysis – the state at which practical reason cannot tell the agent which action to do – with the tiebreaker to be found in the agent’s non-rational impulses,⁷ the dualism can be extended to accommodate cases in which an agent recognizes that she has both reasons to act morally and rationally, and ends up choosing to act one way rather than the other. For example, suppose her reason to act morally is stronger than her reason to act rationally. In such a case we would say that, while she had reasons to act rationally and morally, her reasons to act morally turned out to be *decisive*.

My project here is to develop what I shall call a *communitarian moralist* account of morality. According to communitarian moralism, moral and rational reasons are not commensurable. The result of this, I argue, is that the domain of rational reasons is agent-centric and the domain of moral reasons is found at the level of communities.⁸ For communitarian moralism, agents have rational, but not moral, reasons to act and communities have moral, but not rational, reasons to bring about certain kinds of social organization.⁹

I begin with a discussion of Sidgwick’s dualism of practical reason between reasons generated from what Sidgwick calls the maxim of rational beneficence and the maxim of prudence. After outlining his account of the dualism, I generalize beyond the cases of utilitarianism and rational egoism that he had in mind. I suggest that the maxim of rational beneficence can be broadened to encompass what I call moral reasons and that the maxim of prudence can be broadened to what I call rational reasons.

⁷ Sidgwick, *The Methods of Ethics*, 248.

⁸ I am intentionally avoiding the agent-neutral/agent-relative reasons distinction because it does not quite fit with what I have in mind. By reasons that are “agent-centric” I just mean reasons that agents can have *simpliciter*. My claim here, then, is that rational, but not moral, reasons are part of agents’ psychologies. So if we were to look for an agent’s moral reasons to act that diverge entirely from those projects that she cares about, we will not find anything.

⁹ Again I think it is worth stressing here that my claim is that communities, but not their members, have moral, but not rational, reasons to act.

The third chapter consists of a discussion of the problem of the overdemandingness of morality. In particular, I argue that the problem arises in the form of a demand to act on one's moral rather than rational reasons when the two kinds of reasons conflict. Morality is demanding because it demands that one give up one's personal projects in favour of the moral project. I go on to discuss Paul Hurley's objection that there is a parallel problem of the overdemandingness of rationality. Instead of an overriding demand on agents to be moral, there is a demand on agents to be rational that arises from one's obligations to act on one's rational, rather than moral, reasons. I end the section by showing that these problems arise from a commitment to the two assumptions that rational and moral reasons are commensurable and that they are generated from the personal and impersonal standpoints respectively. I argue that we cannot both accept that moral and rational reasons can be weighed against each other in practical deliberations and that rational reasons are generated from the agent-relative, personal point of view while moral reasons are generated from the agent-neutral, impersonal standpoint.

Using Thomas Nagel's and Paul Hurley's terminology, agent-relative considerations are those considerations about what agents should ultimately do; they are considerations about certain kinds of conduct. In contrast, agent-neutral considerations are concerns about which are the best states of affairs to be brought about; they are considerations about what should happen.¹⁰ Following Hurley, agent-relative considerations naturally have importance in the personal point of view, that point of view in which the rational agent assesses the value of different courses of action entirely as

¹⁰ I take these distinctions from Paul Hurley, *Beyond Consequentialism* (Oxford: Oxford University Press, 2009), especially chapter 6, and from Thomas Nagel, *The View from Nowhere* (Oxford: Oxford University Press, 1986), 164-5.

they relate to her and her own personal projects and commitments, and agent-neutral considerations have impersonal importance, i.e., importance in the impersonal standpoint from which agents step back from their particular circumstances in order to assess the overall value of states of affairs.

Chapter four discusses how the two assumptions that moral and rational reasons are commensurable and that they are generated from the impersonal and personal standpoints respectively are practically inconsistent with each other. That is, with these two assumptions in play, the result of one's practical deliberations is dependent on the standpoint that one occupies. Occupying a standpoint involves taking a certain kind of consideration to be important and ignoring other kinds of considerations. I argue that moral reasons lose their force when they are brought into the personal standpoint and rational reasons lose their force when they are brought into the impersonal standpoint.¹¹ That a particular kind of reason loses its force just means that the weight it carries in a practical deliberation is severely lessened, perhaps in some cases to the point that it has no weight at all; that kind of reason is just not as important in one's deliberation. Because of this, the agent's practical deliberations are stacked in favour of those reasons that are generated in the standpoint that the agent happens to occupy.

In chapter five I argue against Hurley's view that the interpersonal standpoint is an adequate alternative impartial stance to the impersonal point of view. After outlining the arguments for agent-centred permissions and restrictions that motivate Hurley's interpersonal stance, I show that there are at least three problems facing his interpersonal

¹¹ It is important to remember that, while agents do sometimes have personal moral convictions that they very much care about maintaining, those reasons that they have to uphold these particular convictions will be *rational* reasons. In contrast, an agent's *moral* reasons to act are reasons that are insensitive to her personal projects and commitments.

account of impartiality. First I argue that adopting the interpersonal standpoint cannot get us out of the practical inconsistency. Second, I show that while Hurley's account of the interpersonal standpoint can generate agent-centred restrictions against certain impersonally motivated kinds of conduct, it cannot generate any such restrictions against personally motivated kinds of conduct. Finally I suggest that when we find ourselves in the personal point of view, we do not have good reason either to take interpersonal considerations seriously or to adopt the interpersonal standpoint.

In the sixth chapter I consider the possibility that there is no impersonal stance. If we accept that there is no impersonal standpoint and accept that moral reasons, like rational reasons, are generated from the agent's personal point of view, then we can escape the practical inconsistency while maintaining that moral and rational reasons are commensurable. But, as I will argue, this position ultimately fails both because the impersonal stance is entailed by accepting the plausible claim that there can be something that each of the members of a group happen to value and because we cannot give a convincing rationale for affording a special significance to an agent's moral reasons in virtue of their status as moral reasons.

In chapter seven I introduce communitarian moralism as the best way to escape the practical inconsistency discussed in chapter four. I argue that we should reject the claim that moral and rational reasons are commensurable and urge instead that we see moral problems as problems for a community rather than for particular agents. I then use the recurring theme of minimizing interference amongst community members found in the four most prominent moral theories to motivate my discussion of how we ought to go about treating moral problems as a community, rather than as individual agents. I end the

chapter by showing that, on my account, the problem of the overdemandingness of morality that we encountered in chapter three is a non-starter.

Chapter eight consists of a discussion of moral maximization and my communitarian account. I argue that the intuition that agents ought to both maximize the number of moral actions that they perform and maximize the moral content of each action is mistaken. I suggest that moral maximization is not about agents morally maximizing their actions in some way; it is instead about a community's striving to be better. I finish the chapter with my account of how moral criticism of other communities is possible and of what it takes for a community to be considered good or moral.

In the ninth chapter I address the common objection against the impartiality of the impersonal stance that it takes a value to make a value. According to this objection, the impartiality of the impersonal stance is threatened because the values according to which the impersonal stance assesses states of affairs are smuggled into the stance by the people who adopt it. I argue that this objection is not a real problem for me because it mistakes toward whom the impersonal stance is adopted. I suggest that, because the impersonal stance is adopted reflexively – i.e., it is adopted towards itself – it can remain impartial to the particular community, even if it cannot claim to be impartial with respect to other communities.

Chapter ten addresses the problem of moral criticism of agents on a communitarian account. If agents only have rational reasons to act, then one might worry that we cannot morally criticize their conduct. I argue that the moral criticism of agents is possible on my account in several different ways. First I suggest that when agents have an inconsistent set of preferences, some of which are immoral, we can rationally criticize

their conduct and convince them to change their preference set. I go on to argue that in cases in which agents cannot be rationally criticized for holding immoral preferences we can still teach them to cultivate moral preferences, eventually morally rehabilitating them. I conclude the section by showing that the agents for whom moral education and rational criticism fail to remedy their problematic preferences are equally lost on a standard individualistic account of morality.

CHAPTER 2: THE DUALISM OF PRACTICAL REASON

Sidgwick concludes *The Methods of Ethics* by showing that there is a contradiction between the utilitarian maxim of rational beneficence and the egoistic maxim of prudence. He shows that the seemingly self-evident principles that one ought to maximize overall happiness and that one ought to maximize one's own happiness are mutually incompatible. The former is the self-evident principle upon which utilitarianism (and consequentialism more generally) is built and the latter is the self-evident principle upon which egoism sits.¹² His proof of the inconsistency of these two principles has since been referred to as the dualism of practical reason.¹³ Insofar as he thought that utilitarian and egoistic deliberations fall within the scope of practical reason, he thought that there were some situations in which the results of practical deliberations advise the agent to both do and not do some action because it would satisfy one of these principles but not the other. He envisaged a conflict between reasons within rational deliberation whereby no one kind of reason inevitably wins out. Thus for Sidgwick one's egoistic reasons for acting compete against and do not always submit to one's utilitarian reasons for acting.

We can move past Sidgwick's talk of the conflict between utilitarianism and egoism and in its place we can see that the dualism of practical reason more broadly picks out the internal conflict between moral reasons on the one hand and rational reasons on the other. The internal conflict that I have in mind here is the opposition of one's moral

¹² There is some controversy in the literature over whether Sidgwick has rational egoism or ethical egoism in mind. David Brink argues that this turns on whether Sidgwick is a reasons internalist or externalist. I pause only to note this debate since its outcome will not directly affect this project. For the sake of my project here, I shall interpret Sidgwick to be discussing rational egoism.

¹³ See David Brink, "Sidgwick's Dualism of Practical Reason," *Australasian Journal of Philosophy* 66, no. 1 (1988): 291-307; and Francesco Orsi, "The Dualism of Practical Reason: Some Interpretations and Responses," *Etica & Politica / Ethics & Politics* X, no. 2 (2008): 19-41.

and rational reasons within the deliberative stance of an agent.¹⁴ When an agent deliberates, she weighs the different kinds of reasons against each other in order to determine which is decisive. Practical deliberation constitutively aims at this internal conflict's resolution. Recall that moral reasons for action are those reasons to act that an agent has entirely in virtue of the demands imposed on her by morality and that rational reasons are those reasons that an agent has to act in virtue of the demands imposed on her by her self-interest. We can justify broadening the scope of the dualism in this way because a fundamental conflict between one's moral and rational reasons remains intact on any account of morality as long as moral demands sometimes require an agent to give up her personal projects and commitments in favour of moral ones. For my purposes, a reason includes whatever particular psychological states are necessary for intentional actions.¹⁵ I exclude impulsive action from those actions that are guided by reasons and include those mental states (emotions, desires, beliefs, etc.) that are required for one to be motivated to act from one's deliberations. The space of deliberative action broadly

¹⁴ Ultimately, as we will see, I end up denying that such an internal conflict can happen. But this is not to say that agents will not have internal conflicts anymore; rather such conflicts will naturally occur only among her competing rational preferences. These conflicts become much less mysterious because the statuses of the combatants are just 'rational preferences'. Once we see that moral reasons are not found in agents, the conflicts that remain in practical deliberations are just conflicts amongst agents' rational preferences.

¹⁵ I want to refrain from committing myself to one side or the other in the debate about whether internalism or externalism is true of moral reasons. It might seem like I am committing myself to internalism by including the motivational component in what a reason is. However, I am not advancing a thesis about the metaphysics of reasons here; I am merely defining the term 'reason' in such a way that it can satisfy both an internalist and an externalist. If one is an externalist, then by a 'reason' I mean a reason according to externalism coupled with whatever additional motivational component is required to ultimately result in the agent acting. If one is an internalist, then I just mean a reason according to internalism. My point is that, whatever we call it, there is agreement that certain psychological states are bound up with practical deliberation. If someone objects to me that, e.g., reasons are independent of agents' psychological states, and then shows problems with my account by substituting their meaning of 'reason' for my own, he would have to think, erroneously, that I am referring to 'reason,' whatever definition it has, and not 'reason' as I have defined it.

construed – what Sidgwick called practical reason – is the space where this conflict between moral and rational reasons is found.

One might worry that according to this definition of a reason, if an agent develops a personal project to act on a moral reason, the conflict between moral and rational reasons will go away. But this worry does not get off the ground for two reasons. First, insofar as we might think that an agent's personal moral projects have greater significance than her rational ones, the conflict will re-emerge as a serious deliberative problem about when we ought to give up our personal rational project for our personal moral project. Second, insofar as we do not think this, a conflict between moral and rational reasons will persist, but it will be the same kind of conflict that we have about which of our personal rational projects we ought to participate in at a given time. We will return to these themes in greater detail later.

Of course there will be some agents for whom acting morally just is acting rationally. This is not a reason to deny the dualism, however. Rather, in this case we might say that such agents' moral reasons prescribe the same actions as their rational reasons. But intuitively, cases like this are deviant rather than typical.¹⁶ They can be seen

¹⁶ In "A Reconciliation Project," Gregory Kavka seems to argue for the antithesis of my claim here. He thinks that the alignment of agents' moral and rational reasons to act is more common than their divergence. But Kavka's claim is different from mine in an important way: he thinks that agents often have reason to be moral or have good moral reasons to act, and these reasons are typically decisive for most agents. Or rather, he thinks that agents have good rational reasons to conform to the demands of morality in the sense that their life will go better on the whole should they sometimes sacrifice opportunities to indulge in particular personal projects and instead comply with moral demands. This amounts to agents having some particularly strong consideration – e.g., the prudential preference to have one's life go better on the whole – that will override other considerations – e.g., considerations about in which specific projects the agent wants to participate. If Kavka is right, then we might say that moral reasons often align with such prudential preferences. But we cannot say that on the whole it is typically rational to be moral because we would have to ignore the myriad preferences that agents have outside of their particular prudential preferences. See Gregory Kavka, "A Reconciliation Project," in *Morality, Reason and Truth: New Essays on the Foundations of Ethics*, eds. D. Copp and D. Zimmerman (Totowa: Rowman and Allanheld: , 1984), 101-113. We will return to why we cannot privilege one rational project over all others in chapters six, seven and eight.

as instances in which it just so happens that moral and rational reasons align. And because this alignment of reasons is merely contingent, it is not a counterexample to the dualism of practical reason as I have described it.

CHAPTER 3: THE OVERDEMANDINGNESS OF MORALITY

The dualism of practical reason has at least one problematic consequence for moral theories generally, namely, the so-called problem of the overdemandingness of morality. For our purposes, we will concern ourselves with the demandingness problem as it arises for consequentialism.¹⁷ This familiar problem comes about because consequentialism requires an agent to always do that action of those available to her that will result in the best outcome, where the best outcome is defined by the particular theory of the good that is wedded to consequentialism. (For our purposes we shall say it is human welfare.) Further, on standard conceptions of consequentialism, the agent is morally prohibited from doing any other than the best action. But since there is so much avoidable suffering in the world and agents who are affluent relative to those suffering have the means to help prevent it, the action that maximizes outcomes will be the one that exercises those means. The result is that, within a consequentialist framework, the moral demands for any agent who finds herself able to help will require her to give up some or perhaps many of her own personal projects and commitments in order to maximize human welfare. Insofar as the agent's pursuit of these projects and commitments best maximizes her self-interest, she has rational obligations to continue to pursue them. But the demands of morality prohibit her from fulfilling her rational obligations since fulfilling them will not maximize human welfare.

¹⁷ While I go on to discuss consequentialism specifically, there is good reason to think that the demandingness problem can arise within other moral theories to a greater or lesser extent. For examples of how the problem can come about within a deontological or contractarian framework see Garrett Cullity, "Demandingness and Arguments from Presupposition," in *The Problem of Moral Demandingness: New Philosophical Essays*, ed. Timothy Chappell (Basingstoke: Palgrave Macmillan, 2009), especially sections 1.3 and 1.4. For an example of the demandingness problem within virtue ethics, see chapter four of Lisa Tessman, *Burdened Virtues: Virtue Ethics for Liberatory Struggles* (Oxford: Oxford University Press, 2005).

Recall that those reasons that one has to act in whatever ways maximize one's self-interest are rational reasons and that those reasons one has to abide by the demands of morality are moral reasons. The overdemandingness problem arises from the dualism of practical reason in the following way. Consequentialism requires agents to act only on their moral reasons, and rationality requires agents to act only on their rational ones.¹⁸ Sidgwick thought that this just shows that there are two irreconcilable routes to take in one's practical deliberations and that we cannot give reasons to prefer one route to the other. But this by itself does not yield the overdemandingness problem. The problem develops because of how one acquires one's moral and rational reasons respectively.

Rational reasons are acquired from within the personal standpoint – that point of view from which the agent is concerned with her own circumstances and how they affect her life for better or worse. Deliberations about which action to do from this stance will take into account the details of the agent's immediate environment, her values, projects, commitments, prudential considerations and also her mental states and character. The action that is rational to do will be the one that best maximizes her welfare, however it affects the welfare of others.

On the other hand, moral reasons are generated from within the impersonal standpoint – that point of view in which the agent steps back from her own personal circumstances and assesses the circumstances of humanity as a whole. That is, moral reasons are those reasons that take into account the relevant impartial features of the impersonal standpoint, rather than only the particular features of an agent's personal point

¹⁸ The actions that maximize human welfare and a particular agent's welfare are not always mutually exclusive. It is certainly possible for there to be overlap between the two. But even if there is overlap, the basis of the requirements remains different, and, insofar as there are parts that do not overlap, the possibility of conflict between the two persists.

of view. Deliberations about what to do from this standpoint will take into account the aggregate welfare of humanity. The action that it is moral to do for any particular human agent is the one, of those available to her, that best maximizes overall human welfare.

We should not find talk of standpoints too mysterious. Distinguishing between different standpoints in practical deliberations amounts to little more than taking some specific kinds of considerations (e.g., considerations that do not bias some agents over others) seriously and dismissing the strength of others (e.g., considerations that privilege some agent(s) over others). For example, while I might have strong personal reasons to spend my money traveling, those reasons cannot be taken seriously from the impersonal point of view because they violate the impartiality of the impersonal. Put differently, there are better ways for me to spend my money in order to achieve the best overall states of affairs than to go on vacation, and there is no *prima facie* plausible rationale for me to privilege my personal projects over the projects of others. From the impersonal point of view, each agent's projects bear the same weight as each other's and must be sacrificed or pursued only insofar as they conduce to the best overall state of affairs regardless of the agent's personal preferences.

The demandingness problem arises because the actions that are prescribed from reflection on one's impersonal, moral reasons are very often not the same as the actions prescribed by one's personal, rational reasons. Acting on one's moral reasons sometimes requires one to give up one's personal projects and commitments. An agent of relative affluence may have to devote her time and resources to alleviating as much suffering as she possibly can, letting any personal projects she may have fall by the wayside.¹⁹ So if

¹⁹ Though a rational preference to do the moral thing on occasion can motivate her rationally as well as morally, it would be incredibly rare that she is rationally motivated to do the moral thing to the same extent

an agent occupies the impersonal standpoint from which she must take seriously her moral reasons to act, acting in accordance with those reasons is incredibly costly to her. Morality is too demanding because it requires her to give up almost everything that she values personally so that her actions can maximize overall human welfare.

But this problem also goes the other way. In his paper “Does Consequentialism Make Too Many Demands, or None at All?” Paul Hurley argues that consequentialism does not give agents reason to abide by its exacting moral standards. He thinks that because the actions prescribed by one’s impersonal moral reasons diverge so much from the actions one is rationally obligated to carry out, one becomes alienated *from*, rather than *by*, the demands of morality.²⁰ That is, the fact that morality is so demanding gives us good reason not to be moral. We ought not to take our moral reasons seriously, he thinks, because doing so would undermine our rational obligations to pursue our own projects and commitments.

Consequentialism only provides agents with decisive reasons to avoid wrongdoing from the impersonal stance. But since agents naturally find themselves occupying the personal, not the impersonal, standpoint, they will have to put themselves into the impersonal stance if they are going to take moral reasons to act seriously. Of course, doing so requires a *rational* reason to adopt that point of view. And insofar as agents do not usually have decisive rational reasons to adopt the impersonal standpoint for every practical deliberation, they fall short of meeting the demands of morality.²¹

that she is morally motivated to do the moral thing. This level of rational motivation would require that her rational preference to act morally overrides all of her other preferences. We will return later to why it is implausible for an agent to act on only one of her many preferences, even if that is her most highly ranked preference.

²⁰ Paul Hurley, “Does Consequentialism Make Too Many Demands, Or None at All?” *Ethics* 116, no. 4 (2006): 687.

²¹ *Ibid.*, 698.

But perhaps restricting ourselves to talk of standpoints is misleading.²² After all, one's moral and rational reasons are supposed to be commensurable; that is, they are supposed to be of the same kind and are supposed to be able to be weighed against each other. Sidgwick seems to have envisaged the commensurability of moral and rational reasons when he discussed the practical contradiction between utilitarianism and egoism. Once the results come in from one's egoistic and utilitarian calculations, one is required to both do and not do the same action.²³ Sidgwick thought that the outcome of such a circumstance is rational paralysis, only to be decided one way or another by how the agent's non-rational impulses sway her.²⁴ And such a paralysis can only happen if it is possible to look at an action and weigh one's rational reasons to act against one's moral reasons not to act and *vice versa*.

Furthermore, the commensurability of moral and rational reasons is relied on in many thought experiments that are designed to show the demandingness of consequentialism. In Bernard Williams' famous examples of George and Jim, for either to find himself in a real practical dilemma is to presuppose that his moral and rational reasons are competing against each other.²⁵ But if whether one has reason to be moral or rational is determined by the standpoint one finds oneself in, it is not clear why either George or Jim should have difficulty deciding which course of action to take. If Jim finds himself in the impersonal standpoint, he should shoot the prisoner, and if he occupies the

²² I should clarify that it is only misleading if one thinks that moral and rational reasons are commensurable. I discuss the commensurability of the dualism of practical reason presently and we will return to talk of standpoints later.

²³ Sidgwick, *The Methods of Ethics*, 248.

²⁴ Recall that 'rational' for Sidgwick meant 'to be decided by practical deliberation'. I will hereafter refer to this paralysis as practical paralysis in order to avoid confusing Sidgwick's sense of 'rational' with my own.

²⁵ The details of these thought experiments can be found in Bernard Williams, "A Critique of Utilitarianism," in *Utilitarianism: For and Against*, ed. by J. J. C. Smart and Bernard Williams (Cambridge: Cambridge University Press, 1973), 97-9.

personal standpoint, he should not shoot the prisoner. At this point, we might think that talk of standpoints is misleading because it suggests that differing outcomes of practical deliberations from different standpoints cannot be weighed against each other, and this would rule out the possibility of a practical dilemma between rational and moral action. We will return to this in the next chapter.

CHAPTER 4: PRACTICAL INCONSISTENCY

As a matter of fact, agents act on their rational reasons more than they act on their moral reasons. Or put differently, they find themselves in the personal point of view more often than they find themselves in the impersonal point of view. That is why the demands of consequentialism are so onerous. As Liam Murphy points out, if everyone did their part to alleviate easily avoidable suffering in the world, the contributions of individual agents to the consequentialist effort would be relatively small.²⁶ The demands on individual agents are extreme, at least in part, because people often do the rationally self-interested action when they could do the moral one instead.²⁷ Sidgwick was wrong to think that when the constituents of the dualism of practical reason conflict, the result is practical paralysis. Though practical paralysis is certainly possible, it seems that the more prevalent outcome is rational, rather than moral, action.

During an agent's practical deliberation, she might ask herself the question, "Why should I be moral?" And often she will have difficulty finding an answer that will outweigh her rational reasons to refrain from acting morally. The difficulty she has in answering this question is a result of a deep inconsistency between talk of standpoints and the assumption that moral and rational reasons are commensurable. She has difficulty answering this question for this reason: what counts as the right answer to the question "Why should *I* be moral?" must be situated within the personal standpoint and must tell a story about why one's moral reasons to act override one's rational reasons. The question situates the agent in the personal standpoint because it asks why she, not some other

²⁶ Liam Murphy, "The Demands of Beneficence," *Philosophy and Public Affairs* 22, no. 4 (1993): 275.

²⁷ It is, I think, implausible to suppose that this results from a widespread epistemic failing on the part of agents. That is, it is implausible to think that agents typically think, but are mistaken in thinking, that their self-interested actions will yield the best consequences for human welfare.

agent, must be moral. She searches impersonal, moral projects for the same features that pick her out to participate in her personal, rational projects and finds nothing.

In the preceding chapter, I suggested that talk of standpoints is misleading because it might rule out the commensurability of one's competing reasons to act. But talk of standpoints is only misleading inasmuch as one is committed to the commensurability of moral and rational reasons to act. As we will see, this is because deliberating from standpoints is inconsistent with the commensurability of the constituents of the dualism of practical reason. This is not a logical inconsistency, but rather a practical one – one cannot hold that moral reasons are generated from the impersonal standpoint and rational reasons from the personal one and hold that moral and rational reasons compete on equal footing when one asks oneself if one should act morally or rationally. And this practical inconsistency arises in whichever standpoint the agent happens to occupy.

For the consequentialist that we have been discussing, reasons generated from the impersonal, moral point of view are generated by the obligation to fulfill the overall project of maximizing human welfare. Achieving the desired outcome of this project requires the active participation of at least some agents. But the agents who are required to participate are not discriminately required. That is, it is necessary for *some* agents to carry out certain actions that will maximize human welfare, but the project says nothing about *which* agents ought to participate. So, while the requirements on me as an agent who wants his actions to further the aims of the consequentialist project are specific – i.e., there is only one action of those available to me that will best further the aims of the project – it cannot tell me why I should participate in the project to begin with, rather

than someone else. In fact it could be the case that, even while I recognize the importance of the consequentialist project, I decide to act in my own self-interest because I think that others will participate in the moral project.

Of course, one might object that the consequentialist project requires *all* agents to do that action of those available that maximizes human welfare. But it is unclear how a consequentialist account can say this since it is conceivable that the same outcome can be achieved whether or not all agents are acting in accordance with it. If, for example, Murphy is right that total compliance with the consequentialist project will yield a relatively small amount of sacrifice for each agent, it is plausible to think that if I were to forsake my responsibility to comply with the consequentialist project, it can still be realized if everyone else sacrifices a little more. Here we can see the familiar theme of a prisoner's dilemma emerge: each agent does best if she defects and all others comply. But in order for me to do better by defecting, it must be the case that the same project can be fulfilled whether I participate in it or not. Put another way, I can only do better by defecting if I can enjoy the same benefits from the fulfillment of the project that I could have if I were to participate.

So what is at stake in the project is the overall outcome, not who participates in the project and how it is brought about. I might exempt myself from the moral project on rational grounds because one less person behaving morally is not going to jeopardize the overall project. And even if it is true that all agents ought to participate in the moral project, it is not clear why an agent's obligation to participate in the collective moral project should outweigh her obligation to participate in her self-interested projects. At

best the consequentialist project is just one among many that agents have obligations to pursue.

On the other hand, reasons that are generated from the personal, rational point of view are generated by the obligation to fulfill the personal projects of the agent, whatever they might be. The fulfillment of such projects requires the active participation of the agent. Insofar as the projects are *hers* and not someone else's, if she values them at all, she cannot shirk her responsibility to bring them about.²⁸ For example, part of my valuing my own project of having a successful career consists in my active participation in advancing that career. I must often do that action of those available to me that best advances my career. If my career were to further itself, whether I act or not, the result would be unsatisfying since my project is not only to have a successful career, but also to have it, to the greatest extent possible, *by my own efforts*.²⁹ Thus there is an asymmetry between the moral project of maximizing human welfare and my project of having a successful career. In the second case, my active participation is *necessary*, while in the first case it is not. The impersonal moral project of realizing maximum human welfare can be fulfilled even if I choose never to participate in it. But the converse is false: my personal project of having a successful career cannot be fulfilled if I choose never to participate in it.

²⁸ Of course there are cases in which agents do not have to participate in their personal projects in order for the projects to be realized. But while these cases are possible, there are a multitude of others that *do* require the agent's active participation. Likewise, of course, there are some moral projects that do require the agent's participation. We will return to a discussion of these projects later. For the moment it is enough that many moral projects do not require the agent's active participation.

²⁹ It could not be said that my aspirations for becoming an astronaut were answered if I was selected by a lottery that picks citizens at random to become astronauts on the next mission to the international space station, forgoing the rigorous training usually required. In this case we might say that I am not *really* an astronaut, or that I have not *really* achieved my dream. Rather, I am at best a passenger on the space shuttle and at worst a safety risk because of my failure to have acquired the knowledge and skills needed to be an astronaut.

What's more, if I have a personal stake in my active participation in the moral project, then I have a rational reason, not a moral one, to participate in it. That is, I have a rational reason to take seriously and act on my moral reasons. But even in this case I do not have rational reasons to always and everywhere participate in the moral project. Rather, I have one self-interested project among many to participate in the moral project. The moral project itself remains silent about *which* agents choose to participate in it, as long as enough participate for it to be realized.

But the case in which an agent occupying the personal standpoint chooses to act on her moral reasons because doing so is part of the particular personal project that she wishes to fulfill is one of convergence between moral and rational aims. The practical inconsistency between the commensurability of moral and rational reasons on the one hand, and the talk of standpoints on the other, is better shown in the case where the aims do not converge. The distinction between the personal and impersonal standpoints matters for the agent who asks from the personal point of view whether she should do the moral thing after consulting her moral and rational reasons for action. The question for her is whether she should actively participate in the moral project or in one of her rational projects. But as we have seen, the active participation of the particular agent is not required for the success of the moral project because the moral project picks out *some agent* while her rational projects just pick out her. The result is that, when the agent weighs her preferences for the success of the moral project against her preferences for the success of a personal project (e.g., having actively participated in achieving the success of her career), she will inevitably choose the latter because, unlike the former, it requires her active participation. If she chooses the former, she directly hinders the success of her

personal project, but if she chooses the latter, she only indirectly hinders the success of the moral project. Her choice to do the rational thing does not *automatically* hinder the success of the moral project. But the converse is not true: her choice to do the moral thing automatically impedes the success of her career. Thus, when she asks the question, “Why should *I* be moral?” it is not surprising that she has difficulty giving an answer.

The practical inconsistency arises because moral reasons to maximize human welfare are brought into the personal standpoint from the impersonal standpoint where they carry their weight. Just as the impersonal standpoint ignores personal considerations that discriminately privilege some agents over others, the personal standpoint ignores the force of impersonal considerations. Within the personal standpoint, the effects of an agent’s actions on others only matter inasmuch as those effects will cause others to behave in advantageous or disadvantageous ways toward the agent. Because of this, the force of the impersonal, moral reasons to always do the action that maximizes human welfare will be lost within the personal standpoint.

Of course this runs the other way as well. If we ask the question of whether one has a reason to act rationally from the impersonal standpoint, the same considerations will apply. The force of an agent’s rational reasons to pursue her personal projects will be lost when they are brought into the impersonal standpoint because they are reasons *just for that agent*. They will have weight only inasmuch as the increased welfare that results from the satisfaction of those projects for the agent contributes to the welfare of humanity as a whole. The force of such reasons is generally very small. So, from within the impersonal point of view, it is not plausible to suppose that rational reasons will win out over moral ones.

The incompatibility between talk of standpoints and the commensurability of the constituents of the dualism of practical reason arises in the following way. In order for moral and rational reasons to be weighed against each other they must be both brought into the personal or both brought into the impersonal standpoint. But tearing a reason from the standpoint where it was generated inevitably causes it to lose its force. A practical inconsistency surfaces because the reasons that were generated from within the occupied standpoint will almost always win out over the ones that were generated from the other standpoint. If we want to tell a plausible story about the interaction between moral and rational reasons, we have to either accept that each is generated from a different standpoint, or reject that they are commensurable. If we try to accept both, we find ourselves in the practical inconsistency that I have just outlined whereby the reasons generated from whichever stance an agent as a matter of pure accident happens to be in will always be decisive.

CHAPTER 5: HURLEY'S INTERPERSONAL STANCE

I have argued that there is a deep practical inconsistency embedded in the application of the dualism of practical reason to moral theory. When one's impersonal, moral reasons conflict with one's personal, rational reasons, the outcome of one's practical deliberation will be greatly influenced by the particular standpoint that one happens to occupy. This consequence is most commonly seen in the so-called problem of the overdemandingness of morality. The problem arises because the practical deliberations of agents who occupy the impersonal standpoint will usually result in the agent doing the moral thing at the expense of her personal projects and commitments. But of course the parallel problem of the overdemandingness of rationality can be shown to arise from the other standpoint. If the agent happens to occupy the personal standpoint for her practical deliberations, she will usually choose to do the rational thing at the expense of the moral project.

Some writers have suggested that this only shows that the impersonal stance cannot sufficiently provide agents with reasons to act morally. As the default impartial stance that is adopted toward agent conduct, it gives agent-neutral reasons for one to act. But many philosophers think that good moral conduct cannot be adequately explained by just agent-neutral reasons. In "A Critique of Utilitarianism," Bernard Williams famously argued that utilitarianism poses a real threat to an agent's integrity – i.e., it threatens the agent's pursuit of those projects and commitments that give her life structure and in part constitute her moral identity³⁰ – because it gives her only agent-neutral reasons to maximize the best overall outcomes. He thought that such a result is morally problematic

³⁰ There is some controversy in the literature over exactly how to interpret what Williams has in mind by an agent's integrity. I take this definition from Williams, "A Critique of Utilitarianism," 116-7. For a detailed discussion of the different views see Greg Scherkoske, "Whither Integrity II: Integrity and Impartial Morality," *Philosophy Compass* 8, no. 1 (2013), 40-52.

because the impersonal stance is merely a higher-order tool of assessment of lower-order projects. A traditional act utilitarianism adopts the impersonal stance toward the personal projects that directly result in agents' pleasure or pain, or their happiness or suffering. Those personal projects are evaluated from the impersonal stance, and the outcome determines which projects the agent is morally required to pursue and which she is not. But on such a model, the impersonal stance adds nothing to what it evaluates. The moral significance of the outcome of impersonal evaluation is completely dependent on the pursuit of the personal projects of individual agents. Williams concludes that the utilitarian cannot get her account up and running without recognizing the moral importance of the personal projects of each agent. That is, utilitarianism necessarily presupposes the moral importance of the personal point of view.

Samuel Scheffler uses Williams' remarks on the moral importance of the personal point of view to motivate the case for agent-centred permissions within a hybrid consequentialist framework. He thinks that the moral importance of the personal point of view can be captured by allowing the agent to give more weight to her own projects than to the impersonal consequentialist project.³¹ The result, he argues, is an agent-centred prerogative that allows for a good consequentialist to sometimes not maximize overall outcomes in favour of maximizing her own ends. But while we can make a good case for agent-centred permissions, Scheffler thinks agent-centred restrictions are *prima facie* implausible on the grounds that "it is very hard to explain how it can be rational to forbid the performance of a morally objectionable action that would have the effect of

³¹ Samuel Scheffler, "Prerogatives without Restrictions," *Philosophical Perspectives* 6, Ethics (1992), 377-397.

minimizing the total number of comparably objectionable actions that were performed, and would have no other relevant effects.”³²

Recently, Hurley has argued that Scheffler’s agent-centred permissions actually provide all the materials needed to make a case for agent-centred restrictions.³³ Hurley thought that if we accept both that the personal point of view has moral importance and that, as a result, we can provide a plausible rationale for agent-centred permissions, then we must also accept that the same rationale provides agent-centred restrictions. This is because agent-centred permissions are generated when each agent recognizes the moral importance of her own personal point of view. But an agent cannot consistently recognize that her personal point of view has moral importance for her while holding that another agent’s personal point of view does not have moral importance for him. The thought is that, insofar as I can see that my projects and commitments have moral importance for me, I cannot deny that your projects and commitments have the same moral importance for you. Once I accept that the projects of each agent have moral importance for each I have to recognize that, since I am sometimes permitted to not bring about the best overall outcome, they are sometimes permitted to do the same. And since they are permitted to do the same, I am restricted from interfering with the morally permissible pursuit of their personal projects; I cannot force them to maximize the overall good when they are permitted not to. Thus we can generate agent-centred restrictions through the recognition of the moral importance of the personal point of view for each agent by each agent.

Recognition of agent-centred permissions and restrictions results in a different kind of impartial reason than those generated from the impersonal point of view. Such

³² Scheffler, “Prerogatives Without Restrictions,” 387.

³³ The following paragraph summarizes Hurley’s argument for agent-centred restrictions in chapter six of Hurley, *Beyond Consequentialism* (Oxford: Oxford University Press, 2009).

impartial, agent-relative reasons have been called deontological reasons by Thomas Nagel,³⁴ second personal reasons by Stephen Darwall³⁵ and interpersonal reasons by Hurley.³⁶ These reasons are generated by the imperative for each agent to recognize that every other agent's personal projects bear the same moral significance as she takes her own to have. They are, therefore, impartial because they attribute the same moral weight to the personal projects of each agent. When they are taken seriously, they will give agents impartial moral prescriptions about which actions ought to be done, rather than impartial, impersonal moral prescriptions about which states of affairs ought to be brought about.³⁷

Hurley thinks that this can motivate deliberation from the interpersonal standpoint as an impartial alternative to the impersonal standpoint.³⁸ Adopting the interpersonal standpoint involves recognizing the equal moral significance of each agent's projects for them; and the actions that are prescribed by the impartial agent-relative reasons that are generated will be the ones that do not interfere with the permissible pursuit of the projects of others. Such a standpoint finds a middle ground between the exacting demands of both the impersonal and personal standpoints. We can thereby temper the demands both of rationality and morality by adopting the interpersonal stance during our practical deliberations.

But while this alternative impartial standpoint might seem like a plausible candidate for practical deliberation, it has at least three problems. First, it does not resolve

³⁴ Nagel, *The View from Nowhere* (Oxford: Oxford University Press, 1986), 165.

³⁵ Stephen Darwall, *The Second-Person Standpoint* (Cambridge: Harvard University Press, 2006), 7.

³⁶ Hurley, *Beyond Consequentialism* (Oxford: Oxford University Press, 2009), 152.

³⁷ *Ibid.*

³⁸ In fact he thinks that the impersonal standpoint is derivative of the interpersonal standpoint, but he introduces the interpersonal standpoint as an alternative to rather than the grounds of the impersonal.

the tension within the dualism of practical reason. Instead, adopting the interpersonal stance just places boundaries around the courses of action available to the agent, prohibiting her from acting on the most extreme personally and impersonally generated reasons. She therefore cannot kill someone even if doing so would save the lives of many more, nor can she torture someone even if she were to be richly rewarded for doing so. But once these restrictions are respected, all other actions are available to her, whether they are personally or impersonally motivated. The conflict between moral and rational reasons continues on, the only difference is that it takes place in a smaller arena. The practical inconsistency persists once all of the relevant interpersonal reasons for acting or abstaining are taken into account.

Second, the impartiality that Hurley claims is enjoyed by the interpersonal stance mistakes the extent to which an agent is entitled to pursue her projects without interference in virtue of the moral permissibility of her actions. Remember that agent-centred restrictions are generated by the recognition of the same moral permissibility for others to pursue their own projects that I have to pursue my own projects. This is because, insofar as my actions are morally permissible, I am entitled to act without interference from others. But this only seems to make sense if others try to justify interfering with my actions on moral grounds, not as interfering *simpliciter*. Hurley seems to admit as much when he says that “only if others take into account the independent moral significance of my personal point of view for them will my indirect enslavement to the impersonal standpoint be prevented.”³⁹ We can see that recognizing the moral significance of others’ personal points of view will give us restrictions about what sorts

³⁹ Hurley, *Beyond Consequentialism*, 159.

of actions we can require them to perform on impersonal grounds – they are permitted to further their personal projects if they choose.

But what Hurley leaves out is why we should think that interpersonal considerations should prevent others from being enslaved by the personal standpoint. Since agents' entitlement to non-interference seems to apply only to the impersonal point of view, it leaves me open to interfere with the projects of someone else in order to further my own ends. My personal point of view has independent moral significance only insofar as my pursuing my projects is worthwhile for me. Its moral value is derived from the value of my projects *for me*, independent of the value of their outcome from the impersonal point of view. This is because the moral value of the personal point of view is calculated independent from the brute impersonal value of doing such projects. The happiness I derive from doing something that I enjoy is accounted for in the impersonal calculus. But if we afford the personal point of view additional moral significance, that moral significance will be over and above its value in the impersonal calculus. The additional moral value of the personal point of view is a function of how my participation in my projects and my upholding my commitments is valuable to me.

As a result, the means that I take in the pursuit of these outcomes (1) only contribute to the independent moral value of my personal point of view inasmuch as they uphold or violate my *personal* convictions about doing or not doing certain kinds of actions and (2) only have impersonal worth inasmuch as they contribute to or impede the realization of the best overall outcomes. But if I happen to not have personal convictions against coercing or harming others, and since when I exercise my prerogative and pursue my personal projects instead of impersonal ones I ignore the impersonal value that those

projects have, it seems that, on Hurley's account, I am permitted to coerce and harm others in order to pursue my personal projects. Because others are only *morally* entitled to non-interference with their projects, I am not required to refrain from interfering when I am pursuing my essentially non-moral, personal projects.

Finally, even if Hurley accepts that the dualism persists within the interpersonal framework and even if he can incorporate agent-centred restrictions against the enslavement of the personal point of view, the interpersonal account of impartiality faces a further problem. Recall that Hurley argued that agents are alienated from, rather than by, a consequentialist morality because they do not usually have good reason in the personal standpoint to adopt the impersonal stance. The same objection can be levied against taking interpersonal considerations seriously. Sure enough, when we find ourselves in the interpersonal stance we have good impartial, agent-relative reasons to comply with prohibitions on certain kinds of actions. But what is still required is a reason to adopt the interpersonal stance from the personal point of view. If Hurley is right that we can generate agent-centred restrictions from the interpersonal stance, we have no reason to take them seriously when deliberating from the personal point of view, nor do we have good personal reasons to always adopt the interpersonal stance for our practical deliberations.

Without a plausible alternative impartial standpoint, we are once again left with the practical inconsistency between the assumption that moral and rational reasons are commensurable and the assumption that such reasons are generated from the impersonal and personal standpoints respectively. As we saw in chapter four, if these reasons are commensurable then they cannot compete on an even footing because of their disparate

origins. It seems that in order to escape the practical inconsistency we must either reject the origin of these reasons for action or we must reject that they can be weighed against each other.

CHAPTER 6: DENYING THE IMPERSONAL

The first strategy we might attempt to employ to get out of the practical inconsistency is to deny that moral reasons are generated from the impersonal point of view. On this account, moral reasons and rational reasons are both generated from within the personal standpoint. They are commensurable because they are each kinds of reasons for agents to act. In this case, the dualism of practical reason is situated within the agent. In her practical deliberations she will weigh her moral reasons against her rational reasons and then act on whichever comes out on top. On this view, moral reasons are distinguished from rational reasons based on their status and their origin. We might say that one's moral reasons reflect one's moral projects and one's moral projects are some among other non-moral projects that it is rational for one to pursue. On such a view moral and rational projects can be weighed against each other. But moral projects have a special status in virtue of the kind of project they are, and the corresponding reasons for action have a special status in practical deliberation in virtue of being reasons to participate in essentially *moral* projects.⁴⁰ Moral reasons and projects have this standing because of the considerations from which they are generated, namely considerations about the objective good of the outcomes that flow from acting on them, about the objective good of the actions that they prescribe, or even of objective value of pursuing moral ends as such. We will return to why affording such a status to moral reasons and projects is implausible later in this chapter.

⁴⁰ On this view agents will not have rational reasons to participate in moral projects. Rather, they will participate in moral projects because of such projects' special significance as *moral* projects. Agents' participation in moral projects cannot be exhaustively explained by their rational reasons to act. Instead it would have to be understood as a consequence of moral demands that non-impartially pick out particular individual agents to participate in some specified moral projects.

But the position that moral and rational reasons are both personally generated is untenable for another reason. In order to escape the practical inconsistency, it is not enough to find that some moral reasons and projects are generated from the personal point of view. The stronger claim that *all* moral projects are generated from the personal standpoint would have to be made. That is, getting out of the practical inconsistency while maintaining that moral and rational reasons are commensurable requires us to reject that there is such a thing as the impersonal stance. It requires denying that there are moral considerations about a group of agents that are insensitive to the particular projects of its individual members. Although this sort of denial might be controversial, it is not untenable *prima facie*. It is, however, incompatible with the claim that some things, actions or projects are objectively valuable. By objectively valuable, I mean that it is valuable to all members of a specified group – in this case a group of moral agents.⁴¹ Just as something that has subjective value to an agent is in her interest to pursue, something that has objective value to a group is in that group's interest to pursue. All that is required, in this case, is that the group is a group of agents, and any further features of each of its members are irrelevant to the judgment that the group's pursuit of some objective value is in its interest. The impersonal stance falls out of the notion of objective value because it just is the stance from which only the impartial interest of a group has value. The position that moral reasons and projects are generated exclusively from the personal stance is untenable because it requires both that some things are objectively valuable, which supposedly justifies their special statuses in practical deliberation, and that there is no such thing as the impersonal standpoint. But as we have seen, the

⁴¹ To be clear, I mean valuable *to*, not valuable *for*, the agents.

impersonal standpoint is entailed by the acceptance of objective value and so one cannot without contradiction get out of the practical inconsistency by taking this view.

Once again one might look to Hurley's interpersonal stance as an alternative impartial standpoint to the impersonal stance. We might think that, with another impartial evaluative stance at our disposal, we can do without the impersonal one. We saw before that there were three problems with the interpersonal stance. First, adopting the interpersonal stance does not solve the practical inconsistency, it merely restricts the domain in which impersonal and personal considerations can conflict. Second, the agent-centred restrictions that it prescribes can only be restrictions on certain kinds of impersonal conduct; the interpersonal stance cannot give restrictions about how one pursues one's personal projects. Third, we saw that adopting the interpersonal stance from the personal one involved elusive personal reasons similar to the ones that would motivate us to adopt the impersonal stance; Hurley still cannot give us good reason to take interpersonal considerations seriously from the personal standpoint, nor can he give us a good reason to adopt the interpersonal standpoint from the personal one.

For our current discussion, the first objection can be accommodated by accepting that, once all interpersonal considerations are taken into account, all that is left are personal reasons for acting. We can also grant, optimistically, that if Hurley's account is sufficiently refined, he can respond to the second problem. Finally, if, for the sake of argument, we accept his claim that the personal stance is in some sense derived from the interpersonal stance, the third objection is no longer a problem. But even now the interpersonal stance cannot be an adequate candidate for an alternative impartial stance to the impersonal standpoint.

Hurley argues that the interpersonal standpoint arises as the impartial stance motivated by the moral significance of non-impersonal reasons to act. Intrinsic moral goods are those things that have moral worth independent of the weight that they carry in an impersonal calculus. If we can attribute intrinsic moral worth to something like rights without having recourse to the impersonal stance, we get the interpersonal stance in which the reasons bound up with the intrinsic moral worth of rights have the potential to be decisive. Those interpersonal reasons that have the potential to be decisive are reasons for or against certain kinds of conduct, not certain states of affairs. In contrast, objective goods are those goods that are common currency amongst agents, goods such as pleasure or happiness. Further, objective goods are features of states of affairs. The evaluation of the presence of objective goods in a state of affairs is the operation of the impersonal stance. It is objective goods, not intrinsic ones, that are the inputs to the impersonal calculus.⁴²

The impersonal stance just is the assessment of the maximization of – or, less strongly, the having more of rather than less of – objective goods. While it might be possible to recognize the value of intrinsic goods, by itself this cannot sufficiently satisfy the intuition that it is better to have more goods than less goods. Recognizing the value of intrinsic goods in a non-impersonal stance only involves recognizing that they have value. The further step of understanding that we ought to pursue that value because it will result in a better state of affairs is a matter of taking rights to have objective value (and in that way they will become features of states of affairs) and then adopting the impersonal

⁴² Intrinsic goods can be taken into account in the impersonal calculus, but only insofar as they conduce to higher amounts of objective goods. For instance, we can admit rights only insofar as their protection will promote more happiness, pleasure, etc. than their violation. Intrinsic goods can only be cashed out in terms of objective goods in the impersonal standpoint.

stance toward them. Prohibitions against certain kinds of conduct that violate rights fall out of the fact that one ought to act in ways that conduce to realizing things of intrinsic value. But as soon as we add that it is better to respect rights more than less, we are comparing the value of states of affairs in which the degree to which rights are respected differs. The problem is that, now that we are comparing states of affairs and not kinds of conduct, we find ourselves in the impersonal stance again.

Since we are not permitted to have the impersonal stance on the alternative escape from the practical inconsistency, we must also relinquish the intuition that it is better to have more, rather than less, of a feature of value in a state of affairs. But while it might be the case that we are sometimes not required to maximize the goods that are common to all because of considerations from the interpersonal stance, we have to take the further step and jettison the intuition that having more of a good is better than having less, applied relative to a group. But once we get rid of this intuition, what we are left with no longer appears plausible. This is because we will have to accept that while within the personal stance a rational agent has, at least most of the time, a rational obligation to himself to choose actions that will conduce to his greater overall pleasure/happiness/etc., we as a group do not have a similar obligation to ourselves to realize our greater overall pleasure/happiness/etc. It would be inconsistent to accept the claim that it is better to have more, rather than less, personal good brought about by discharging my obligation to act rationally and then reject the claim that it is better to have more, rather than less, group good brought about by us discharging our obligation to act together in the group's interest. The obligations only differ in the number of agents involved, which is hardly grounds for thinking that, for the group, there is no obligation at all. Just as I have an

obligation to myself to realize my best overall state of affairs, it seems that we have an obligation to ourselves to realize our best overall state of affairs. The absurd consequence that, while I can evaluate states of affairs for their value to me and my projects, we cannot evaluate states of affairs for their value to us and our projects falls out of any view that rejects the impersonal standpoint, even if that view offers an impartial alternative like Hurley's interpersonal stance.

CHAPTER 7: COMMUNITARIAN MORALISM

The failure of this approach suggests that in order to rescue ourselves from the practical inconsistency, we should consider rejecting the commensurability of moral and rational reasons. We should reject the notion that within practical deliberations one weighs one's moral reasons against one's rational reasons and then acts on whichever wins out. Put another way, we should reject the thought that during one's practical deliberation, one has two different kinds of reasons – moral and rational – and one must choose which of these reasons will be decisive. We should no longer think of an agent's deliberations as involving her asking herself the two questions: “Why should I be moral?” and “Why should I be rational?” with the answer to the former taking the form of moral reasons to act, and that of the latter taking the form of rational reasons. Instead I will argue that we should think of the agent asking the same questions, but answering them both with rational reasons for action.

As we have seen in the classical tradition, when the moral and rational actions open to the agent are different, she will have difficulty weighing these reasons against each other. After all, she is asking why *she* should be moral and why *she* should be rational, as if the moral and rational projects are both hers and hers alone. It is here that the mistake lies. The classical tradition makes the assumption that both moral and rational reasons are individualistic when, as far as I can see, they are not. While rational reasons surely are individualistic, it is not clear why we should think that moral reasons are individualistic as well. We will come back to this point later, but what is important is that when an agent asks herself the questions “Why should I be moral?” and “Why should I be rational?” the reasons that she is searching for are individualistic. They are reasons

that pick her out and only her. For the majority of her personal projects, she can answer why she ought to act rationally easily: she ought to act rationally because the rational reasons that she gives herself are essentially individualistic.⁴³ A reason to pursue my personal projects, to maximize my utility, to serve my wellbeing, is only a reason for *me*. These reasons are special because they do not and cannot pick out any other agent to bring about my projects, utility, or wellbeing.

Rational reasons are individualistic because it is impossible for me to have a reason for someone else to act rationally, complete with all the relevant phenomenology of having a rational reason as such. This is because, as I have said above, what I have in mind is more than the mere belief that doing X will conduce to my wellbeing – such a belief can be shared with other agents. But on my account a rational reason is the complete kit of the relevant beliefs, aims, motivations, emotions, etc. that are required for rational action. They remain individualistic even though someone might have reasons to help me satisfy my rational ends. It would be foolish to say that my rational reasons are also automatically hers. We would be committed to saying much more than that she participates in the fulfillment of my projects; we would have to say that she must share the very same bundle of mental states as me, something that is, of course, impossible. At the very least, parts of that bundle are emotions that cannot be shared between us because for me they are self-regarding and for her they are other-regarding.⁴⁴

⁴³ By ‘individualistic’ I not only mean that they are reasons that are indexed to particular agents, but also reasons that are bound up with agents’ characters, projects, interests, etc. We should also not confuse individualistic reasons with traditional agent-relative reasons. Agent-relative reasons can be shared between different agents inasmuch as their circumstances are alike. But because I am working with a theory-neutral idea of reasons, complete with all the important psychological states, I mean to reject that agents’ rational reasons can be shared. I use the term ‘individualistic’ rather than ‘agent-relative’ mostly in order to avoid confusion about differing notions of what a reason is.

⁴⁴ Take, for example, the difference between my fear of failure and her fear for my failure.

This is not to say that there are no other-regarding rational reasons for action. But such reasons are reasons for an agent to act that happen to involve the satisfaction of another's ends as part of the agent's reasons. Other-regarding rational reasons like my reason to participate in your projects are still reasons *for me*, not you, to participate in those projects. Though there might be some overlap between our reasons to participate – perhaps we share the same beliefs about the actions that conduce to realizing the project and we surely share the same aims (i.e., the project itself) – there are important ways in which they are distinct. My motivations for helping will reflect my character, commitments and other facts about me and your motivations for acting will reflect respective things about you.

But while rational reasons are individualistic, moral reasons are not. We can see why by looking at the demands that morality imposes on agents. These demands are not individualistic because moral considerations are those considerations about the right sorts of human intercourse that minimize interference with the pursuit of agents' respective projects. Any formulation of the moral project has the peaceful coexistence of self-directed humans living in a community as its telos.⁴⁵ That is, a moral project will tell a story about how the humans within a moral community should be able to pursue their self-directed ends with small risk of interference from others. What counts as a good moral community just is one in which the overall flourishing of its members is high because its members are able to realize most of their projects most of the time. A moral project attempts to maximize the pursuit of as many individual projects as it can and it

⁴⁵ I use a very loose sense of 'community' that can be expanded to encompass the world as our ability to make moral change around the world increases. Further, one might worry that different conceptions of human flourishing – for instance those that do not fall out of western liberal values – will result in a different definition of a good community. I return to this discussion later in chapter nine.

restricts the pursuit of projects that will undermine the possibility of pursuing other individual projects. But I am not committed to the claim that any project whose pursuit interferes with the pursuit of other projects must be restricted. Rather, I have in mind a threshold at which the participation in projects that take away the possibility of pursuing other permissible projects ought to be restricted. Where this threshold lies is admittedly a grey area, but we need not be exact here. My view allows for some interference between agents pursuing their ends, as long as that interference does not fall on the wrong side of this threshold.

We can see this minimization of interference between the pursuits of agents' projects within normative theory as it stands. The consequentialist project has as its telos the best overall state of affairs for human good, whatever the good is, because such a state of affairs would be one of minimized human harm.⁴⁶ Moreover, it is a state of affairs in which members of a moral community coexist in such a way that relieves harm from and conduces to the good of others. Thus, consequentialism aims at a minimized level of interference in the projects of others, insofar as that interference would be harmful to them and impede the realization of their ends.

The Kantian project has the minimization of interference of ends built into its categorical imperative. On the formula of universal law, Kant says: "Act only on that maxim through which you can at the same time will that it should become a universal law"⁴⁷ In this form, the categorical imperative charges agents to deliberate about which

⁴⁶ While it is sometimes contested that the consequentialist minimizes harm rather than maximizes good, I do not think that is a problem for my view. It might mean that, depending on one's interpretation of the telos of consequentialism, the threshold for the restriction of projects lies in a different part of the grey area that I mentioned above. This, I think, is an advantage to being un-exact about what counts as too much interference in the projects of others.

⁴⁷ Immanuel Kant, *Groundwork for the Metaphysics of Morals*, ed. Jonathan Bennett (2005) accessed August 9, 2013, <http://www.earlymoderntexts.com/pdf/kantgrou.pdf>, 24.

maxims to act on and to only act on those maxims that all agents can act on without interfering with the actions of others. This rules out numerous areas of human conduct that hinder the realization of the self-directed projects of agents. On the formula of humanity, Kant charges agents to always act in ways that treats humanity as an end in itself.⁴⁸ This calls on agents to always act on those maxims that do not take other agents to be mere means to a given agent's ends. The formula of humanity prohibits one from acting on maxims, acting on which would have one interfering with others in their pursuit of their ends because such maxims ignore that others even have their own ends. Further, some maxims, such as those that would prescribe acts of coercion or enslavement of others, not only interfere with the projects of the coerced and the enslaved, they also foster the resentment of the coerced and enslaved toward their oppressor. Communities in which coercion and slavery are rampant are examples of bad social organization because coercion and slavery severely impede the pursuit of the self-directed projects of the coerced and enslaved. And on top of this, the resulting resentments will undermine future cooperation between oppressed and oppressor parties.⁴⁹ Such social tension not only impacts the flourishing of the oppressed agents, but also the flourishing of the oppressors insofar as some projects require the consensual participation of all parties involved in their pursuit.

Virtue ethics commands agents to act as the virtuous person would act. It requires agents to cultivate the virtuous character traits. Perhaps unsurprisingly, the virtues that one ought to cultivate are those traits that will conduce to greater peaceful cooperation

⁴⁸ Immanuel Kant, "The Metaphysics of Morals," in *Practical Philosophy*, ed. Mary J. Gregor, trans. Mary J. Gregor (Cambridge: Cambridge University Press, 1996), 516-7.

⁴⁹ For a detailed discussion of the effects of such resentments see Katie Stockdale, "Collective Resentment," *Social Theory and Practice* 39, no. 3 (2013), 501-521.

amongst humans. For instance, in *Nicomachean Ethics*, the disposition that Aristotle finds between the vices of obsequiousness and contentiousness marks the virtuous mean of what is sometimes called *friendliness*, though Aristotle himself left it unnamed. Such a person, Aristotle said,

seems to be concerned with the pleasures and pains of social life; and wherever it is not honourable, or is harmful, for him to contribute pleasure, he will refuse, and will choose rather to give pain; also if his acquiescence in another's action would bring disgrace, and that in a high degree, or injury, on another, while his opposition brings a little pain, he will not acquiesce but will decline.⁵⁰

Friendliness is a virtue because it finds the mean that conduces to the best kind of social life. What's more, Aristotle admits that the virtue of magnificence is actually a vice of excess. But magnificence is a virtuous trait because, despite its technical status as a vice, cultivating it is "neither harmful to one's neighbor, nor very unseemly."⁵¹ Friendliness is a virtue that, when cultivated, will give agents good reason to restrict their interference in the lives of others – in his example, one who has friendliness will choose to inflict a small pain on his neighbor so that he does not have to severely injure or disgrace the other. Magnificence is justified as a virtue just because its pursuit does not interfere with others in harmful ways.

A contractarian account of morality prescribes moral demands that will minimize interference between agents as well. Thomas Hobbes motivates his social contract by juxtaposing the anarchist state of nature, where life is "solitary, poor, nasty, brutish and short,"⁵² with the state under a sovereign in which there are constraints upon actions that

⁵⁰ Aristotle, "Nicomachean Ethics: Book IV," <http://classics.mit.edu/Aristotle/nicomachaen.4.iv.html> (accessed 08/08, 2013). Ch. 6.

⁵¹ *Ibid.*, Ch. 2.

⁵² Thomas Hobbes, *Leviathan*, <http://oregonstate.edu/instruct/phl302/texts/hobbes/leviathan-c.html#CHAPTERXIII> (accessed 08/09, 2013), Ch. XIII.

interfere with others' lives, the intuition being that such constraints are justified because nobody would want to live in the state of nature. On his contractarian account, T. M. Scanlon argues that agents ought to act on principles that they can reasonably expect others to be in agreement about, the idea being that nobody will agree to principles that fundamentally conflict with most of their projects.⁵³

We can see from this discussion that morality more generally demands that agents act together in ways that conduce to a minimization of interference in the pursuit of self-directed projects. Or put another way, morality imposes demands on human communities. It demands that those communities be communities that foster a level of peaceful coexistence of their members that makes it possible for them to pursue most of their projects most of the time.

We can now understand this solution to the practical inconsistency as a dispute over who is the target of moral demands. The dispute is between the classical moralist who thinks that morality makes demands on agents and what I will call the communitarian moralist who thinks that morality makes demands on communities.

When the classical moralist asks what sorts of things the moral agent ought to do she might imagine the agent asking herself the question "Why should I be moral?" As we have seen, this question can be difficult to answer, especially when it is slightly altered: "Why should I be moral, rather than rational?" But the difficulty she has answering this question might be indicative of the fact that it is a bad question. I suggested this earlier in my discussion of the practical inconsistency within the dualism of practical reason. Recall that the answer to this question, what I call the *agent-moral question*, too often eludes us

⁵³ T. M. Scanlon, *What we Owe to each Other* (Cambridge: The Belknap Press of Harvard University Press, 1998), 4.

because, in asking it, we are bringing impartial considerations into the personal standpoint.

In contrast, the communitarian moralist recognizes these difficulties and asks a different question. Rather than “Why should I be moral?” the communitarian moralist asks “Why should we (as a community) be moral?” And the answer to this question comes easily. It is the justification for the telos of normative theory: a society that fosters peaceful coexistence amongst its members is a society in which its members best flourish. The communitarian will tell a story about how the conditions for the flourishing of the members of her community are set by having the proper institutions upheld within the community. Moral questions asked from this perspective can be easily answered by citing facts about the flourishing of individual agents. And moral demands are demands for a community, not for an agent; they are demands that call for the community to adopt the right institutions, whether or not each and every member of the community decides to abide by them.

What I have in mind when I talk about human agents flourishing in a community is different from the standard use of flourishing in a eudemonistic virtue ethic. Human agents flourish, on my account, if and only if they are able to pursue most of their projects most of the time. The more projects they can pursue and the more often they are able to pursue them, the more they will flourish. Thus, when I say that a community can implement institutions that best promote the flourishing of its members I mean that the institutions that it can justify implementing will have to allow for the pursuit of as many of its members’ individual projects as possible most of the time. This comes with a qualification. I want to include the thought that some projects cannot conduce to human

flourishing in a community because they essentially interfere with the projects of others. Rape, murder and torture projects are some examples of these severely harm-inducing projects, and no doubt there are many others. So a community that implements institutions that lead to the ability for its agents to flourish more will allow its members to pursue as many of their personal projects as they can as often as they can as long as those projects are not essentially harm-inducing.

But in addition, agents might have to be restricted from pursuing other projects on the grounds that they happen to interfere with the projects of others too much, even if such projects are not essentially interfering. It might be the case that the flourishing of the members of a community will be greatest when certain projects that are not essentially harm-inducing are restricted. Determining how to have the greatest human flourishing is the goal that we strive for when answering moral questions. On a communitarian account, we can think of the community as a rational agent itself that only wants the flourishing of its members and so implements various social and political institutions and conventions that will best maximize the flourishing of its members. I should note that this is merely an analogy and should not be taken literally, but it is a good way to phrase how a community will try to address moral questions and problems. We can make sense of which institutions a community should adopt by imagining what is the best way it can go about realizing its goal of having the greatest flourishing of its members.⁵⁴

⁵⁴ We should also be careful that we do not take some actions that a community takes as a means to a state of social organization as actions that it is instrumentally *rational* for the community to do. We might instead introduce talk of instrumentally *moral* actions for a community, which would simply be those actions that a community takes as a means to its members' flourishing. But these details are beyond the scope of this project. It is important at this point only that we do not mistakenly attribute instrumentally rational actions to communities.

This communitarian account, I think, has many advantages over the classical one. First and foremost, the classical account allows for the practical inconsistency to be resolved only if we jettison the impersonal standpoint. But we have seen that such a position is incoherent if we hold some sort of account of objective value because any account of objective value will entail the impersonal standpoint. Furthermore the classical account has difficulty answering the agent-moral question, “Why should I be moral rather than rational?” because the agent’s projects and commitments inevitably get in the way of her moral convictions. But the communitarian stance resolves the practical inconsistency by rejecting the thesis that moral and rational reasons are commensurable with each other. In fact, the communitarian rejects the notion that agents have moral reasons at all; moral reasons are not possessed by individuals, but by a community. The result is that the agent-moral question does not make sense for the communitarian. She never has to answer the difficult question “Why should I be moral rather than rational?” she need only answer the more easily answered question, “Why should we (as a community) be moral?”

I do not mean to imply that a community is a kind of agent in its own right, capable of having beliefs, hopes, desires, etc. Rather, what it means for a community to possess moral reasons, for example reasons to adopt a particular institution like a public school system, is for a subset of its members to deliberate on behalf of their community, taking into account only moral reasons.⁵⁵ Those members who deliberate on behalf of the community choose institutions using only impartial considerations about the flourishing

⁵⁵ It matters little who actually deliberates on behalf of the community. Most often it is the community officials and politicians that do so, but sometimes it might begin with members who do not hold positions of power. What matters is that such deliberations *really are* on behalf of the community and are not self-serving for deliberators. When these deliberations end up admitting various self-serving features, they are open to criticism from others in the community. It is, after all, a human activity and is therefore subject to the possibility of mistakes, biases, etc. and can therefore be revised and done better. We will return to this discussion in a little more detail in the following chapter.

of its members. They adopt the impersonal stance towards the community about which institutions best conduce to human flourishing and they will then implement those institutions. And they are not subject to Hurley's problem that they lack motivation to go from their personal point of view to the impersonal stance because those who perform these deliberations will be motivated by personal projects to bring about a better community – this is, for example, why some people become policy makers: they often have the betterment of their community as a personal project.

On my communitarian account, agents can still have moral projects and they can still do the moral thing. But agents having moral projects and acting morally is not something different from agents having self-interested projects and acting on them. On a communitarian account, a community's projects are moral projects that it has moral reasons to participate in. In contrast, an agent's 'moral' projects are just those personal projects that she has rational reason to pursue that are such that they either uphold the community's moral institutions or directly resist violating them. For example, suppose that an agent has a personal project to keep promises. This project is just one of his personal, rational projects that he has good reason to pursue and it is only called a *moral* project because of his community's moral institution that promise keeping should be promoted. Hereafter I will maintain this distinction between a community's moral projects and agents' moral projects.

So, an agent's moral projects are just some individual projects among others, and they are personal, rather than impersonal. Of course, even if we accept that an agent's moral projects are just a particular kind of personal project that she has good reason to pursue, the intuition might persist that, while her moral projects are motivated self-

interestedly, she still ought to prefer the satisfaction of her moral projects to the satisfaction of her non-moral projects. That is, it might be thought that while an agent's moral projects are just a particular kind of her self-interested project, those moral projects should be afforded a certain special significance over and above her other projects. The reason behind this is the idea that moral requirements ought to override one's other self-interested projects. And the rationale for the claim that moral considerations ought to be overriding lies in a story about the why moral projects are importantly different from rational projects.

This story comes in many different colours. On a consequentialist model, moral projects promote the overall good of humanity, whereas self-interested ones promote the good of only one human. On a Kantian model, moral projects protect the intrinsic worth of human agency rather than merely giving the satisfaction of one agent's inclination. On a virtue ethical account, moral projects direct the agent to activities that conduce to her own well-being instead of activities that consist in the agent indulging in her vices. And finally, within a contractarian framework moral projects lead to good relations amongst agents rather than to situations where the desires of only one agent are satisfied and interpersonal relations are harmed.

These differences supposedly justify the importance of moral projects over self-interested ones because the goals or effects of moral projects are arguably more valuable than those of self-interested projects. For consequentialism, the good of all is necessarily greater than the good of one. On a Kantian framework, respecting the intrinsic worth of human reason is immeasurably more valuable than the satisfaction of any number of

human inclinations.⁵⁶ Virtue ethics explicitly denies the worth of pursuing vice at the expense of virtue; what it means to become virtuous just is both to pursue one's own wellbeing and to cultivate the virtues while not succumbing to vice. Contractarianism relies on the rationale that it is better for everyone considered individually if they live in an environment in which everyone participates in, or at least behaves in ways concordant with the mutual advancement of, each other's projects. If good interpersonal relationships are absent in a community, then the projects of individual agents that require the help of others will be constantly frustrated.

If it is indeed right that the outcomes or the teloi of the moral projects are more valuable than those of the self-interested, non-moral projects, then one could plausibly argue that moral projects should be afforded special significance in practical deliberations. But in order for this to work the value of moral projects and the value of non-moral projects must be of the same kind. That is, they must each be valuable in exactly the same way; the moral project just has more of the same value than the non-moral project. But for whom are these respective projects valuable? If we are to give moral projects special significance in the practical deliberations of the agent, it must be the case that they are, as a matter of fact, more valuable *for the agent*. That is, the special significance can only be plausibly justified if we find that moral projects are more valuable than non-moral projects for the agent who has them.

But the story that each of the prominent moral theories has for why moral projects are more important than non-moral ones finds that moral projects are more valuable

⁵⁶ Kant refers to the "good will" (having the appropriate moral dispositions and acting on them) as something that is unconditionally good, implying that all other goods are only conditionally good. He goes on to suggest that because of this, the good will should override mere desires. Kant, *Groundwork for the Metaphysics of Morals*, 5-7.

because they result in better forms of social intercourse in a community. That is, moral projects are more valuable because their realization affects all moral agents in a community, whereas the realization of non-moral projects affects very few moral agents in a community (or perhaps only one such agent). And, so the story goes, we can attribute more value to moral projects by the sheer weight of numbers.

The communitarian is on board with the justification for why moral projects have great value for a community. But having great value for a community is not sufficient to justify the claim that moral projects ought to be afforded special significance in individual agents' practical deliberations. The story that consequentialism, Kantianism, virtue theory and contractarianism each give in order to justify giving moral projects this sort of special significance mistakes for whom moral projects are valuable. They make a mistake in thinking that moral projects ought to be more valuable for the individual agent because the realization of a moral project is more valuable for the community than the realization of any of an agent's non-moral projects. But without justifying the claim that it is more important *for a particular agent* to satisfy her moral projects instead of her non-moral projects, the special significance cannot be given to her moral projects. The value of moral projects for a community was never in question on a communitarian account, but it is a mistake to suppose that a moral project's value for a community should be the same as its value for an individual agent.

This is not to say that the justification for the significance of moral projects over non-moral ones is completely unimportant for the communitarian account. The problem that we have been discussing lies in the assumption that this significance ought to be found in the agent's deliberations. But the stories that proponents of each of the

prominent normative theories will tell about why moral projects have this significance is exactly what justifies the community's adopting institutions that better promote the flourishing of its members over ones that do not. These stories are important because they explain why moral projects have greater value – not in the individual stance, but from the standpoint of a community. One of the strengths of my communitarian account is that the importance of moral projects for individual agents does not require and cannot be explained by an appeal to impersonal value. An agent has moral projects only because of the value that they have for her, whereas a community has moral projects because of the value they have for all of its members. The importance of moral projects to an agent is measured in exactly the same way that the importance of her non-moral projects is, namely, only with respect to her set of personal preferences. While impersonal value determines what institutions her community ought to implement, it is silent about which project the agent ought to choose to act on.⁵⁷

Moral problems, then, are problems for a community.⁵⁸ On my view, their solutions will not be different kinds of agent conduct. Instead they will have to do with finding the right sort of social organization on the political level and the right sort of moral conventions on the social level. Since moral questions are asked and answered from the community's point of view, the implementation of the answers can be straightforward movements for social and political change. Because we have rejected the assumption that moral and rational reasons are commensurable, the particular projects of individual members of the community will have no weight when the community answers

⁵⁷ We should not be too worried about the diversity of values within communities. Inasmuch as there is disagreement within a community about which are the appropriate values that persist in its standing institutions there is room for criticism and reform. We will return to this topic in chapters eight and nine.

⁵⁸ Just like agents' moral projects are not really *moral* projects, agents' moral problems are not really *moral* problems. They are, rather, rational problems, the details of which are discussed in chapter ten.

to the demands of moral projects. And the converse is true as well: the particular moral projects of the community have no weight on the agent choosing amongst her ends.

Let us unpack these two claims. In the first case, the community will account for the general kinds of preferences of its members when it answers moral questions. It takes the general preference for its members to pursue their personal projects, whatever they happen to be, under consideration when it prescribes which institutions should be implemented and which conventions should be upheld. It also takes general facts about human flourishing into account when answering these questions. Each of these considerations will both help it to say which are the right sorts of institutions and help it to rank the level of moral importance of the different aspects of human flourishing and general preferences of its members. For instance, we might plausibly think that the ability for agents to pursue their personal projects is morally worthy, but that the fact that human flourishing typically requires living a longer, rather than shorter, life is overriding. The result of this might be a form of social organization that resembles a welfare state, whereby agents' projects can be justifiably limited if doing so will allow agents to typically live longer lives.

But the deliberations of the community will be insensitive to the particular content of the projects of its constituent members. The adoption of institutions that are justified from particular agent-centred preferences is not a moral affair – or if it is, it is a bad one. For example, if the institution of condemning those who dislike golf were to be implemented, it could only be justified because the contingent preferences of golfers were mistaken as preferences for an intrinsically good activity, the participation in which conduces to human flourishing. What might be required instead is that the community

fosters institutions that do not impede human flourishing. This gives its members latitude to choose between the sorts of activities that they pursue.

In the second case, the moral mandate of the community does not put pressure on its constituent agents to always and everywhere be moral. Nor does it attempt to foster the preference for its agents to act morally. It is silent about the conduct of the agent. Morality is not a set of “Thou shalt...” rules imposed on agents, it is rather a collection of cohesive institutions within which agents are free to pursue their ends, whatever they might be.

Recall that it does not follow from this that there is no such thing as an agent acting morally. But agents who do the moral thing are better described as acting without breaching moral institutions rather than as acting morally. When an agent abstains from murder or rescues a drowning child, he does not do so from an overpowering moral demand to forsake his personal project of attaining vengeance or saving his suit. Instead he acts only from his own preferences. Perhaps he prefers to not bloody his hands, or he prefers that the innocent do not suffer. His so-called moral action is not moral because he obeyed generalized imperatives that overpowered his other projects. It is moral because of the circumstances, because it upholds the moral institutions of his community.

Of course there are other actions that are performed and motivated by the fact that they are moral. These are actions that directly buttress or even improve the community’s moral institutions. But despite the fact that the motivation for agents to perform these actions is strictly because they are moral, they do not reflect a moral imperative for agents to bring them about. Again, these actions are evidence only of the agent’s preference to act morally. The agent’s contingent preference for her community to have

stronger, better institutions is sufficient for her to act so as to better realize them. Because of this, the actions that she chooses will reflect what *she* values, not what is objectively valuable. For example, a human rights activist might dedicate her life to the project of leveling arbitrary inequalities amongst members of her community. Such an activist might believe that what she is doing ought to be done, morally speaking. And she might also experience deep satisfaction at her project's successes and frustration at its failures. But the *reason* that she chooses to be moral by fighting for equal rights cannot be found within an imperative to maximize on the pure overarching moral good for all. The reason that she is a human rights activist is first and foremost because *she* values human rights above other goods. And she values human rights whether or not they are, in fact, an intrinsic moral good.

So the moral actions of agents are not going to be in conflict with the agent acting for her own interest. Rather, these actions will be actions that affect the moral institutions of her community. She is motivated to do the actions either because they are the means to the satisfaction of one of her preferences, because she values the flourishing of her community or because the performance of the action has intrinsic value to her. These reasons for doing the moral thing are explained not by an imperative imposed on her by the demands of the standards of ethical conduct, but only by her rational reasons to act.

On my communitarian account of morality, the dualism of practical reason does not sit within the agent's deliberations. It is rather a dualism between reasons for a community (what I have called 'moral reasons') on the one hand, and reasons for an agent (what I have called 'rational reasons') on the other. These two kinds of reasons are incommensurable – they cannot compete against one another. For the communitarian, the

conflict that an agent experiences between doing the moral thing at the expense of one of her non-moral projects is just a conflict about which of her self-interested preferences she will try to satisfy. Her reasons to do the moral thing are reasons that motivate her to bring about one of her projects. But crucially, it is only one of her projects among others.⁵⁹

When we phrase the broad problem of overdemandingness in communitarian terms, we can see that it is no problem at all. Recall that many moral theories, especially consequentialism, have been accused of being too demanding for agents because they require agents to give up many of their personal projects for moral ones. If an agent is to be truly moral on these theories, she must *always* act morally.⁶⁰

On the communitarian model, this problem is a non-starter because the circumstances in which it arises are implausible. If we translate the problem into communitarian terms, we have an imperative for the agent to act only on her preferences to do the moral thing at the expense of her other preferences, whatever importance their objects have to her. That is, we have to suppose that an agent must arbitrarily ignore all but one of her projects. That she has to maximally advance only one of her projects is arbitrary because, as I have argued above, she cannot rely on the status of the reasons she has to act on that project. Since the communitarian rejects that these particular reasons have special authority in virtue of being *moral* reasons, she cannot even plausibly describe the circumstances in which the problem of overdemandingness arises. The

⁵⁹ Paul Hurley has argued for a similar point in *Beyond Consequentialism*. He brings it up as a criticism of consequentialism. Since consequentialist demands can only be met by agents who have consequentialist aims as one of their projects, it is merely one project among many and thus lacks the ability to demand that agents fulfill the consequentialist project to its full extent. But this does not count as an objection to my view since I am not interested in justifying the overriding rational authority of morality like a classical consequentialist might.

⁶⁰ In "Moral Saints," Susan Wolf argues that strict adherence to these theories requires agents to be "as morally good as possible." See Susan Wolf, "Moral Saints," *The Journal of Philosophy* LXXIX, No 8 (1982): 421.

problem arises on a communitarian account only if it is plausible for an agent to maximally advance just one of her projects, and, if this is actually plausible, only if it can then be shown why it is not arbitrary that she maximize her participation in her community's moral projects and not some other, personal project of hers.

CHAPTER 8: COMMUNITARIAN MORALISM AND MORAL

MAXIMIZATION

One of the consequences of my communitarian account is that it absolves agents of the obligation to morally maximize their conduct. For a communitarian, the agent is neither required to always and everywhere do the moral thing, nor is she required, when she chooses to act morally, to maximize the moral content of her action – she need not do the ‘best’ moral action of those available to her. The former sense of moral maximization I call *quantitative* moral maximization because it is the maximization of moral acts. When an agent quantitatively maximizes moral actions, she does as many moral actions as she can, or strives to make all of her actions moral actions. The latter sense I call *qualitative* moral maximization because it is the maximization of the moral worth of particular actions. For an agent to qualitatively maximize her moral conduct, she must choose that action that is the ‘best’ moral action of those available to her. The communitarian stance denies that an agent has obligations to morally maximize either quantitatively or qualitatively. As a result of this, doing the moral thing is no longer horrendously exacting on moral agents. On my view a moral action for an agent is just a self-interested action that intersects in some way with her community’s moral institutions. Since agents are not required to morally maximize, they are free to act morally when and how they choose.⁶¹

⁶¹ In her paper “Morality as a System of Hypothetical Imperatives,” Philippa Foot argues that moral demands must be hypothetical, not categorical, imperatives. That is, she thinks that agents take moral actions as a means to some other end, rather than taking them to be an end in themselves. Foot’s view differs from my own in that she thinks that taking moral actions to certain ends is morally praiseworthy, but to others it is not. (314) For example, when one acts honestly as a means to being an honest person it is praiseworthy, but if one does so in order to be recognized as being honest by others it is not. On my account, such actions are not necessarily praiseworthy or blameworthy because of the ends for which they are carried out, rather they are nothing more than merely rational actions that are praiseworthy or blameworthy only insofar as they intersect with the community’s institutions. What’s more, as we will see in chapter ten, there is nothing automatically wrong with agents acting morally in order to, e.g., be seen to be moral. Philippa Foot, “Morality as a System of Hypothetical Imperatives,” *The Philosophical Review* 81, no. 3 (1972), 305-316.

The sense that one ought to engage in the maximization of moral actions is driven by the intuition that, while doing the moral thing once is good, doing it a second time is better, a third time even better, and so on. This intuition is, I think, the right one, but it has been wrongly placed in the agent's deliberations. When we think of moral maximization, we think of maximizing the moral actions of agents. The thought is that if we can get agents to repeatedly discharge their moral duties at the expense of their non-moral duties, we can better realize moral projects. However, as I have argued above, this way of thinking wrongly assumes that the conduct of agents is the focal point of moral projects. Because sustaining the good institutions of a community, rather than the conduct of individual agents, constitutes genuine moral projects, it would be a mistake to locate moral maximization at the level of individual conduct. Instead, we should morally maximize at the level of the community's institutions.⁶²

It does not make sense to quantitatively maximize agents' moral conduct for exactly the same reason that we could not plausibly formulate the overdemandingness objection on the communitarian model. An agent's maximizing moral conduct just is her only ever acting on one of her projects – namely, her moral project – and never on any others. But this is at best entirely arbitrary and at worst simply implausible because she cannot justify realizing only one of her projects and ignoring all of her other projects on the one hand, nor can she justify consistently choosing her moral project to maximize on the other, unless that project is given special significance. Of course, on a communitarian model, there is no such special significance; her moral project has the same status and importance as her other projects. And since it cannot have a special significance, it cannot be plausibly argued that the agent ought to act on her moral projects at the expense of her

⁶² We will return to what it means for a community to maximize its institutions below.

non-moral projects simply in virtue of their status as *moral* projects. The result is that, from a communitarian stance, moral agents cannot be required to quantitatively maximize moral conduct.

The communitarian also denies that moral agents are required to qualitatively maximize their moral actions. That is, an agent does not have an obligation to choose the best moral action of those available to her when she chooses to act morally. This is because if she had the obligation to qualitatively maximize, the action that would discharge it would be completely insensitive to her particular projects and character. Her course of action would be determined with reference only to the considerations of the moral project in which she has decided to participate. It would be determined only based on what is the best thing to do, morally speaking. Or put another way, the agent's action would be chosen for her based on what it is best to do *for that moral project* (determined by impersonal value because it is moral) rather than by her based on what it is best to do *for the agent – her* – who chose to participate.

This is problematic both because it over-prescribes the action that agents are required to do in order to discharge their obligation to qualitatively maximize and because it ignores important features of what is required for an agent to act. When agents choose to participate in the same moral project, they will each be prescribed the same 'best' action, call it Φ .⁶³ But this only makes sense on a model of moral obligations according to which one has obligations to Φ entirely in virtue of being a moral agent. That is, it only works on a framework in which agents are indiscriminately called upon to Φ in order to fulfill the moral project. This results in the following two implausible

⁶³ Or perhaps the moral project will require a combination of different actions in order to achieve the best overall outcome. But who is selected to do which action will be entirely arbitrary, ignoring the important particular features of individual agents in order to retain impartiality.

claims: 1) only agents' Φ -ing can realize their moral project, and 2) Φ can be carried out simultaneously by many agents. It is not clear why it should be accepted that satisfying a moral project consists in doing one action. Rather, it seems to be overwhelmingly plausible to think that moral projects are best realized through the different actions of many agents working together. And further still, it is not clear why we should think that, in all cases, some moral action can be performed by many agents, regardless of their particular circumstances. The actions of individual agents are constrained by the circumstances in which they find themselves such that 2) is patently false.

While this might seem implausible, it is important because it shows the consequences of the commitment to the view that agents ought to qualitatively maximize their moral actions. The view is implausible because it ignores many of the crucial features of the particular agent who chooses to act morally, features that are added to the equation when she asks herself, "What ought I to do?" When an impartial account of moral obligation requires an agent to do the *best* action, morally speaking, the agent is called on to Φ , regardless of her particular projects and character. Whenever such an impartial story of moral obligation is qualified in some way – perhaps the agent is not required to Φ because she is a trained doctor and can better participate in the moral project by doing some other action – it admits certain non-moral facts into the relevant considerations of what the best moral action for her is. But such facts will be bound up with her personal projects and her character as an agent who wants more than to simply do the moral thing. As a result, the 'best' moral action for a particular agent can never be determined from purely moral considerations. So, an agent cannot be qualitatively required to maximize her moral actions as determined only with respect to her particular

moral project. The ‘best’ moral action for her will be determined partially by the project in which she wishes to participate as well as the nature of her character and her other personal projects.⁶⁴

Defenders of qualitative maximization will argue that their view does not require agents to perform the same *action*; it obliges them to perform certain *kinds* of action. But even so, it is not clear that they can get obligations for kinds of actions, rather than particular actions themselves, from just one’s being a moral agent. This is because an obligation for an agent to bring about certain kinds of actions necessarily admits more features than just her status as a moral agent. It admits the particular context in which she finds herself – she must choose that action of those available to her. And it admits particular features of her character – the action that she must do within the scope of the kind of action that she is obliged to do depends, at least in part, on her ability to successfully bring it about. For example, suppose that either you, a medical surgeon, or I, an accountant, come across a dying man in an alley. For the sake of argument, suppose as well that each of us has an overarching moral obligation to alleviate suffering. Upon inspection we each see that if the bullet lodged in his side is not removed, the man will slowly bleed out. As a seasoned medical surgeon who has treated bullet wounds like this before, you see that your medical expertise will allow you to extract the bullet and treat the wound without a very high risk to the man’s life. As an accountant who knows only rudimentary first aid, I see that I can try to staunch the flow of blood but that I lack the expertise needed to extract the bullet without risking the man’s life. Intuitively, if each of us is to try to help the dying man, what we ought to do is different: you ought to extract the bullet and then treat the wound and I ought to call an ambulance while I attempt to

⁶⁴ Margaret Urban Walker, *Moral Context* (Lanham: Rowman & Littlefield Publishers, 2003), 9-12.

staunch the bleeding. But if the qualitative maximizer is right, then the particular actions that you and I ought to do within those actions that will alleviate the man's suffering will be entirely determined by our mutual status as moral agents. This means that the features about our different backgrounds that determined *how* we ought to go about saving the man's life cannot be admitted into our deliberations since they are things over and above our shared moral agency.

My claim here is not that, e.g., a utilitarian will not admit such particularizing features to her account of which actions a given agent is called on to do in the name of the utilitarian project. My claim is instead that, insofar as these features are bound up with the personal projects, commitments and character of the agent, if a utilitarian insists that they play an important role in determining which action an agent ought to do, the impartiality of the demands placed on the agent is lost. This is because the capacity that one has to contribute to the utilitarian project is bound up with one's personal projects and desires. In the above example, perhaps you only became a doctor because of an interest in anatomy and because you care about helping others. But because your particular interests allowed you to become a doctor, which then gave you your capacity to contribute to the general good, those interests have moral importance that cannot be accounted for in an impartial framework. That is, the utilitarian cannot allow one's capacity to contribute to the utilitarian project to influence which action one ought to do without also admitting the importance of one's particular interests. And insofar as she admits the moral importance of one's interests as well as one's capacity to contribute to the general good, she loses the impartiality of her framework.

The result of this is that the qualitative maximizer cannot enjoy an adequately robust account of which actions it is best for agents to perform unless she rejects that moral obligations are rendered impartially so that she can admit the particularizing features of who is doing the action, where the agent finds herself and how she can best go about doing it.

But we should not be tempted to conclude from this discussion that moral maximization should fall by the wayside. I have argued only for why moral maximization cannot be about agent conduct. But it remains true that both qualitative and quantitative moral maximization are very important from the standpoint in which moral considerations have authority. That is, while it is not true that agents ought always and everywhere do the best moral action, it is true that a community ought always and everywhere strive to have the best social and political institutions. The maximization intuition easily finds its place at the level of the community. On a communitarian model, what it is for a community to be good or moral is for its institutions to best maximize human flourishing. When institutions are implemented in a community, they must be implemented because they will best uphold good human intercourse.

Here we should be careful lest we be misled by the language of maximization. When I say that communities should morally maximize their institutions I do not mean that they ought to reach the state of affairs in which they have all of the maximally good institutions. This way of understanding the goodness of a community – whether the community is a moral one – is both implausible and unhelpful. It is implausible because it requires a complete understanding of which institutions will best conduce to its members' flourishing, and it is unhelpful because when every community fails to even be able to

understand what is required to be a good community, let alone actually become one, the project of becoming a better community appears unrealizable. Instead I want to fall back on the intuition that motivates maximization that I mentioned above: for a community, while adopting a good institution once is good, doing it twice is better, and so on. We should think of a community's moral maximization as its striving to become more moral, given its current institutions, political climate, etc. A good community will be one that has institutions that conduce to its members' flourishing and tries to find and implement better ones. A better community will be one whose institutions help the flourishing of its members more; a worse community is one that helps its members to flourish less.

This might leave open the question of *how* a community ought to go about adopting its institutions. I will not go into the answer to this question in detail here; that is another project in its own right. What little I will say, however, is that such a process will be informed by at least three things: the nature of certain personal projects, facts about what best conduces to human flourishing, and the collaborative effort on the part of the community's members – both those who hold the power and those who do not – to comply with the institutions that are implemented. These features, and perhaps there are more, will conduce to a better community insofar as the community members who uphold them are well-informed, uphold them often, etc. and conduce to a worse community insofar as the community members are not well-informed, seldom uphold them, etc. The ability for a community to be moral depends on the ability of its members to discover and implement the right institutions well. And we should be optimistic about the likelihood that community members sincerely wish to participate in bringing about a better community because, insofar as having the right institutions will set the conditions

for the possibility for the community members to realize their personal projects, each member has a personal stake in maintaining and advancing those institutions.⁶⁵

However, one might worry that having this personal stake in the betterment of the community will not be sufficient to motivate some agents to develop community-bettering projects. Such agents will free-ride on the efforts of others to make a better community. But this is just another case of a prisoner's dilemma. If everyone else participates in community-bettering projects and I do not I will be better off than if I sacrifice the pursuit of some of my projects in favour of community-bettering ones. But the possibility of resolving prisoner's dilemmas is a different, though related, project to this one. An adequately robust account of how we ought to resolve prisoner's dilemmas can therefore be added to my communitarian moralism in order to deal with this objection.⁶⁶

But even if I, as someone who cares about the betterment of his community, find myself in a community in which there are an overwhelming number of free-riders, my project to bring about a better community does not increase, but it might change. There is an important difference between my community here and one in which everyone wishes to participate in community-bettering projects. Whereas the best projects in the community with high participation will likely be those that lead to the adoption of particular institutions, the projects that will serve the community of free-riders best will likely be those that increase the participation of others in community-bettering projects.

⁶⁵ It is important to note that because of this, we do not need to talk about the obligation of the community's members to bring about the betterment of the community. As we will see in chapter ten, inserting the language of impersonal obligation that trickles down from how to have the best community adds nothing to the motivation of its particular members.

⁶⁶ For some plausible candidates see Duncan MacIntosh, "Preference's Progress: Rational Self-Alteration and the Rationality of Morality," *Dialogue: Canadian Philosophical Review* XXX (1991), 3-32. and David Gauthier, *Morals by Agreement* (Oxford: Oxford University Press, 1986).

This only goes to show that the community of free-riders is worse than the community without them. What I must do in each community in order to participate in my community-bettering projects will be different because of the state of the community in which I find myself, but how much I participate in these projects is still determined by how much I care about building a better community.

CHAPTER 9: IT TAKES A VALUE TO MAKE A VALUE

It is controversial within the literature just what adopting the impersonal stance entails. According to the traditional view, adopting the impersonal stance involves becoming the agent of ‘pure’ reason, stepping back from one’s subjective desires, values and attitudes in order to see the world through the lens of pure practical reason. The thought is that from such a stance I can see the ethical truths about the world without my contingent, subjective psychology obscuring my vision. Conceptions of the purely rational agent can be traced back to the ancient Greek deity Apollo. In modern times, the traditional view was championed by Kant whose ethic was founded on the idea that everyone has the capacity to step back from her personal inclinations and desires in order to become an impassive, rational Apollonian saint. In contemporary thought, this view has been taken up by John Rawls who thought that if all rational agents in a society were to step back from the particular features of their lives into what he calls the “original position,” they could all come to agree about how to organize a just society.⁶⁷ In each case, the traditional view involves leaving all of one’s values behind so that one’s deliberations about what it is right to do will not be distorted. Only by deliberating from this standpoint can one see clearly the truth about what one ought to do.

There are at least two standard objections to this view. First, the traditional view has been accused of ignoring the moral relevance of the values that one is required to step back from in order to ethically deliberate well. Feminists and particularists like Margaret Urban Walker have argued that one cannot be required to step back from one’s particular values and commitments in order to come upon the ethical truth of the matter because

⁶⁷ John Rawls, “A Theory of Justice,” in *Reason at Work* 3rd edition, ed. Steven M. Cahn (Orlando: Harcourt Brace & Company, 1996), 264-5.

doing so will ignore particular features of the context that have moral relevance. I think this view is right for reasons discussed in the preceding chapter. As I argued there, this objection gives us good reason to deny that individual agents ought to adopt the impersonal stance in order to be good practical deliberators.

Second, the traditional view has been criticized for violating the principle that it takes a value to make a value. In *Ruling Passions*, Simon Blackburn argues that Kant's pure practical reason and Rawls' original position inexorably exemplify certain agent-centred values. As he says when discussing Rawls' original position:

The real problem is ... the motivational one of commanding respect for what would have been chosen in that position. "What I would have chosen, had I been different in some specified way" is just like "what X would choose". Whether the answer is of any moral interest depends on whether we respect and admire X or whether we think X a broken reed ... Rawls's contractors leave behind ... everything except a dislike of risk, a concern for a fairly long term, and a stripped-down concern for the necessities without which life is bound to be miserable... With only these kinds of concerns ... a person should, Rawls argues, choose a legal and economic system closely resembling those of modern western welfare-state democracies.⁶⁸

Even if we grant that agents are able to assume the original position, whether or not we think the outcome of adopting the position is a good one will be determined by whether the values brought into the stance are respected. There are at least two consequences that are of interest for this project: first, the impersonal stance requires inputs of certain values in order to produce anything practically interesting, something which threatens its claim to strict impartiality; and second, whether the output of adopting the stance is to be trusted is determined by what kinds of values are brought into it.

With this in mind, it is not surprising that the normative component of my account is one that prescribes adopting many of the institutions of a western liberal democracy.

⁶⁸ Simon Blackburn, *Ruling Passions* (Oxford: Oxford University Press, 1998), 273.

The impersonal stance, on my view, is one that embraces the value of the self-determined pursuit of one's own projects with prohibitions against the pursuit of those projects that will severely undermine the self-determined pursuit of the projects of others. The normative theory that I have embedded within my communitarian framework is one that assumes the value of those institutions that secure the self-determined pursuit of its members, and it is through this value that other communities can be morally criticized. This is because, as a western liberal, I have the value of the self-determined pursuit of one's projects built into what it means for humans to flourish. The result is that my communitarian framework will admit a degree of ethical relativism inasmuch as there is disagreement about what it means for humans to flourish.

But we should not be worried by the relativity of human flourishing. In disputes about what makes humans flourish there are some facts that will not be relativized to a particular community. For example, it is not relativized to different communities that humans flourish more if they live longer, eat, drink, sleep and avoid enduring severe pain, all other things being equal.⁶⁹ Furthermore, there is often agreement amongst communities about the relativized facts that conduce to human flourishing. Most communities see murder as something that does not conduce to the flourishing of the victim and often does not conduce to the flourishing of the perpetrator. The extent to which the communitarian framework allows for ethical relativism through disagreement about what counts as human flourishing is therefore quite small. Within the framework, the application of a particular normative theory that is informed by a particular conception of what it takes for humans to flourish may vary in distinct ways. But insofar

⁶⁹ When other things are not equal, this might no longer be the case. If I am being tortured, it might be the case that living longer will not conduce to my flourishing. But the point is that these extreme cases are the extenuating circumstances under which the rule of what it means to flourish admits an exception.

as there is common ground among the conceptions of human flourishing on which the differing theories are built, there is room for each community to morally criticize the other.⁷⁰

But the objection looms that the impersonal stance that each of these communities takes is not strictly impartial because it assumes the value of a particular conception of human flourishing. That is because the occupation of the impersonal stance has a bias built in that favours the community that occupies it over others. The community cannot be strictly impartial when it criticizes other communities – it will often favour its own institutions.

This objection mistakes toward whom the impersonal stance is supposed to be impartial. The impersonal stance is not one that is adopted so that one moral community can assess another on an even playing field. Rather, the impersonal stance is adopted reflexively – toward the community itself – so that it can assess which institutions it ought to adopt without privileging some members and ignoring others. The impersonal standpoint is adopted in order for the community to take some value – the value of human flourishing – and distribute it amongst its members such that most of its members can flourish most of the time. The origin of the value of human flourishing is built into the communitarian framework and the specific features of what it means for humans to flourish will be in part relative to, and in part independent of, the particular community.

⁷⁰ We should not mistake the divergence between communities about what it means for humans to flourish as a matter to be cleared up by one community stumbling upon the truth of the matter and then showing other communities their mistakes. This would only result in each community, having believed they discovered the truth with their particular conception of flourishing, trying to dogmatically teach each other community to see the discovered truth. This would inevitably lead to confusion and frustration about why others stubbornly fail to be persuaded of one's discovery of the truth. Instead we should understand the disagreement among communities as a symptom to be treated, rather than a failure of communities to rightly perceive the world. That is, we should not think of other communities as making a fundamental mistake about the world because the problem to be addressed is the *disagreement* between other communities and our own.

The relativity of human flourishing for a particular community can manifest itself through everything from the contingent beliefs of community officials and policy makers to general trends of etiquette that emerge from the daily rituals of its members. Insofar as there is disagreement over the appropriateness of community institutions and social norms that arise as a result of this relativity, there is opportunity to reform or improve them.

CHAPTER 10: CRITICIZING AGENTS' ACTIONS

On my account, a criticism of an agent's action is, most of the time, a criticism of her rational conduct. Communities are criticized for being immoral, but agents are criticized for being irrational. Yet we want to be able to criticize someone who fails to uphold or directly violates the moral institutions of his community by calling him immoral. And this criticism, we think, ought to really bother him. We can still do this on my account, but when we call him immoral, it is not because he is failing to discharge his overarching obligation to, let's say, abstain from murdering. On my view, we call him immoral because he fails to satisfy his personal project of not becoming a murderer. There is a *rational* failing on his part inasmuch as he cares about not becoming a murderer, but we call him *immoral* just because the project that he fails to realize is a project that intersects with the community's moral institutions that prohibit murder.⁷¹ We can show him that he ought not to kill through an explanation of how becoming a murderer will frustrate many of his other projects, of the external sanctions that he suffers should he choose to kill, of the guilt that he will feel for having become a killer and more. The authority of his obligation not to kill is derived from his preferences, not from the universalization of a principle prohibiting murder.

But sometimes rational agents will not have the right preferences. For instance, it might be the case that an agent, John, does not have personal convictions against murder. In cases like this, it might be objected, we cannot rationally criticize his failure to abstain from killing because he has no anti-murder preferences. Moreover, since the

⁷¹ To be clear, in this case the agent is irrational as well as immoral. He is irrational because he fails to satisfy one of his personal projects that he cares about and he is immoral because that project is one that intersects with his community's institutions that prohibit murder. So calling the agent immoral expresses a judgment that his preference to abstain from killing is being flouted or that he simply fails to have the preference at all, in either case resulting in a breach of his community's anti-murder institutions.

communitarian account denies that there are such things as personal moral reasons derived from overarching moral obligations about individual agent conduct, any criticism that John is behaving immorally will fall on deaf ears. And since he does not happen to have anti-murder preferences, any criticism that he is behaving irrationally is patently false. As a result, the account might seem implausible because it does not seem capable of providing an acceptable account of moral criticism.

The communitarian response to such an accusation will be first that the contingent lack of appropriate preferences is often not an irreparable state, and second that, insofar as my account might suffer from this worry, an alternative, individualistic account of morality will suffer at least as much, if not more. In the first instance, I suggest that moral education will help to remedy the deviant preferences of community agents like John. In addition, I show that the classical moralist can only have claim to the same kind of ability to morally educate John if it admits the same tools that the communitarian account uses; the classical moralist both has to accept, and adds nothing to, the communitarian attempt to remedy John's preferences. In the second instance, I show that the classical moralist will suffer from the same problem if it proves impossible to teach John to change his preferences. I argue that having recourse to overarching moral obligations on individual agent conduct will not force John to change his preferences, whatever they are.

It is easy to see that John's lack of the appropriate anti-murder preferences does not make him a lost cause. It is also clear that John has a serious problem and that his preferences must be rectified – that is why we accuse him of being immoral. But such a problem can often be overcome. For example, we can explain to him how he might be mistaken about some facts that would, if he were aware of them, lead him to give greater

value to the preference not to kill. We might also explain to him why it is in his best interest not to kill, and more broadly why it is in his interest to uphold the community's moral institutions in general. We can explain to him the probabilities of getting caught and the harsh community sanctions on such conduct. We can try to understand his personal history to see why refraining from murder is unimportant to him. We can teach him to embrace empathy, sympathy or other appropriate emotions if he should find himself in circumstances where he would actually kill. We could have, in John's case, exercised good childhood education about the right sorts of behaviour, teaching children to prefer to uphold the community's moral institutions rather than to violate them. And there might be even more strategies that can be used to show John why he ought to, rationally speaking, alter his preferences.

Each of these strategies will be effective in some circumstances and ineffective in others. There will be one that is most reliable and one that is least reliable. The 'right' strategy for moral education will depend on the particular features of John's situation and whichever strategy can be expected to elicit the right change of preferences. For instance, one might think that teaching him to cultivate an intrinsic desire – a desire for something just for its own sake – to uphold the community's moral institutions will prove a more reliable way for him to abstain from killing than teaching him to develop an instrumental desire – a desire for something as a means to the satisfaction of another desire – to abstain from killing as a means to avoiding harsh judicial punishment. And it might be true that the success of the former will have a better impact on his behaviour than the success of the latter. But it might also be the true that the former is less reliably taught than the

latter; it is more difficult to successfully teach an agent to cultivate an intrinsic desire than it is to teach him to develop an instrumental desire.

Some will worry that if he adopts the preference to refrain from killing merely as a means to avoiding harsh judicial punishment, he is still not being moral because he is behaving morally for the wrong reason. The intuition behind this kind of objection is that in some sense John is faking it; he is not *really* moral because he would kill without the deterring institutions in place. This, I think, can be reinterpreted into a claim about the reliability of the strategy of teaching him to adopt an instrumental preference not to kill. We might think that developing an instrumental desire is not enough for John to refrain from murdering because his commitment to holding an instrumental desire is completely dependent on what its satisfaction is a means to. In this case, John's instrumental desire not to kill is dependent on the harsh judicial punishment that he would receive if he were to kill. But if the deterrent is removed, the instrumental desire falls away with it; John was just pretending to be moral, we conclude, he never actually was.

When we see the objection against developing instrumental preferences to uphold the community's moral institutions as a claim about the unreliability that such preferences conduce to John's avoidance of murder, it becomes less problematic. The unreliability of instrumental desires in a context like John's shows two things. First, it shows that developing an instrumental desire will lead John to avoid killing, and second, it shows that, insofar as such desires are unreliable, when John has developed these instrumental desires, his moral education is still incomplete. The unreliability of instrumental desires is only seriously problematic if we think of them as the final stop in moral education. But if we see them as an important stepping-stone in the gradual process

of moral education, they cease to be so alien. The fact that John's desire not to kill is merely instrumental is only evidence that his rehabilitation is incomplete.

But the original problem will resurface when all attempts at moral education fail. There will be some cases in which, while moral education will work for John, it will not convince another agent, Jack, to change his problematic preferences. Suppose that Jack has no anti-murder preferences just like John, but Jack, unlike John, also has highly ranked killing preferences. In fact, Jack finds himself in circumstances in which his killing preferences are so strong that they cannot be overcome by any of his other preferences. In a case like this, both moral criticism and moral education will fall on deaf ears; there is no convincing Jack to desire or act differently.

In Jack's case, the communitarian account will have to give up. There will be no convincing Jack to alter his preferences or to change his conduct. But the classical moralist fares no better. The tools that the classical moralist has to work with here are the same methods of moral education. But the classical moralist will also try to appeal to the overarching or universal moral rightness and wrongness of certain kinds of conduct. Because, in this case, moral education fails, the classical moralist must convince Jack, at the very least, to change his conduct by explaining to him the objective moral wrongness of killing. But it is not obvious why such an explanation would make Jack care that he is doing something morally reprehensible. That is, it is not clear why it should matter *to him* that his conduct is morally horrendous.

We might imagine this case as one in which, were Jack to adopt the impersonal stance in his deliberations, he would have decisive reasons to avoid killing. But Jack is such that he both lacks personal projects that involve adopting the impersonal stance, and

lacks the necessary personal projects to elicit good desire changes. The latter will result in a failure of moral education. And, as Hurley has argued within a classical moralist framework, the former results in Jack's alienation from morality. His personal projects will never get him to the impersonal stance where knowledge of the objective moral wrongness of killing will carry weight in his deliberations.

But the classical moralist will fare worse than the communitarian when it comes to moral criticism. In a case like Jack's, the attempts of the classical and communitarian moralist to criticize Jack's conduct will fail to elicit both a change in conduct and a change in preferences; each fares no better than the other. In a case like John's, however, the classical moralist will require that John recognize the objective wrongness of killing in order for a change of conduct to be a strictly moral one.⁷² Even a consequentialist will want to describe an agent's acting because it will produce the best overall outcomes and an agent's acting because it will conduce to the satisfaction of her personal projects, even if both agents' actions are the ones that bring about the best overall states of affairs. A consequentialist will say, "They both did the right thing, but only one of them did it on purpose."⁷³ The classical moralist does not fare as well as the communitarian in John's case because of the requirement that he recognize the moral wrongness of killing in order for him to escape the criticism of being immoral.

Above I argued that the success of morally educating John is determined by how reliably he abstains from his criticized conduct. I responded to the objection that he

⁷² In fact, Kant famously argued that unless the maxim that directed one's action issued from the categorical imperative, the agent's conduct was not moral. The requirement that I am supposing is not as strong as that; it is simply that the objective moral rightness or wrongness of one's conduct must be recognized – perhaps as well as a personal project, intrinsic desire, etc. to pursue or refrain from that conduct – for the agent to count as more than accidentally acting morally.

⁷³ The fact that the agent either intentionally or unintentionally did the moral thing will not add anything morally relevant to situation for the consequentialist. The point is that whether the agent did it intentionally or not remains an important feature of the context.

would just be pretending to be moral were he to develop an instrumental desire not to kill as a means to avoiding harsh judicial punishment by suggesting that this only shows that his moral education is incomplete. I also suggested that developing an intrinsic desire to avoid killing might be more reliable than developing a similar instrumental desire. But this is only one path among many that could be taken to successfully morally educate John. The communitarian account actually leaves open how one is successfully morally educated as long as the method produces a reliable change in behaviour.⁷⁴ In contrast, the classical moralist runs into difficulty because it requires that all methods of moral education end with some appreciation of the objective moral rightness or wrongness of one's conduct in order for one to intentionally be moral.

Since the preferences that must be cultivated to avoid killing are essentially his, we might have the further worry that, even if he were morally educated by a classical moralist to recognize the good for all of his abstinence from killing, he would not develop the strong commitment to the project of avoiding killing others that we hope for. But in morally educating John, our goal is not to explain to everyone what is best for all; it is instead to explain to *him* why *he* ought to care about not killing others. Our project, then, is to give John a personal stake in the project of not killing others. When it comes to moral education, the strength of the communitarian account over the classical one is that it gives us a way to explain *to John* that he ought not to kill; i.e., it allows us to make him see why *he* ought not to be a killer – why killing is bad *for him*.

⁷⁴ This is not to say that John can be tortured until he is so terrified that he will never kill. That method and those like it will be morally problematic on a communitarian account because they violate the community's moral institutions and should therefore be morally criticized. Under these methods, John's change in conduct and attitudes will typically be unreliable because he might lose faith in the institutions that he is being taught to care about and his change of attitudes and conduct again may be merely instrumental.

CHAPTER 11: CONCLUSION

In this thesis, I have argued that there is a deep practical inconsistency that arises from the acceptance of the two assumptions that rational and moral reasons are commensurable and that rational and moral reasons are generated from the personal and impersonal points of view respectively. As we saw, this inconsistency comes about because of the model of the dualism of practical reason that can be extracted from Sidgwick's remarks on the tension between the maxim of rational beneficence and the maxim of prudence. In order for there to be such a tension, the commensurability of the constituents of the dualism must be assumed. And the assumption about the origins of the constituents comes easily from our intuitions about which features ought to be taken seriously when deliberating about moral or rational projects.

A traditional formulation of the problem of the overdemandingness of morality falls out of the acceptance of these two assumptions. But through accepting these assumptions we also encountered the opposing problem of the overdemandingness of rationality. The practical inconsistency poses a serious problem for deliberating agents because the mere occupation of the personal or impersonal standpoint will require the agent to take only one kind of reason to be decisive. In the personal standpoint, the demands of rationality are so exacting that she is required to ignore the demands of morality, and in the impersonal standpoint, the demands of morality are so exacting that she is required to ignore the demands of rationality. This is a problem because moral and rational reasons can never be weighed against each other on an even footing. Occupying a particular standpoint will cause agents to take one kind of reason seriously simply in virtue of its genesis, rather than by the relevance of its content.

I considered two ways to escape the practical inconsistency and ultimately found both of them to be wanting. First I explored the possibility that Hurley's interpersonal stance can give us a way out of the inconsistency. But using the interpersonal as an impartial alternative to the impersonal is problematic for three reasons: first, it does not resolve but merely constricts the conflict between moral and rational reasons; second, it only places prohibitions on impersonally, not personally, motivated conduct; and finally, it is not clear that agents deliberating from the personal standpoint will have good reason to take interpersonal considerations seriously.

I then discussed the plausibility of denying the origin assumption of the practical inconsistency. Denying this assumption requires that both moral and rational reasons are generated from the personal point of view. But as we saw, we stumble upon at least two problems. First, it is not clear that we can consistently hold a relatively broad account of objective value and deny that we can generate impersonal reasons to act. Second, even if we can somehow justify the claim that moral reasons are generated exclusively from the personal point of view, we cannot supply a plausible rationale for giving such reasons a special significance in our practical deliberations.⁷⁵

It seems to me that the most plausible escape from the practical inconsistency is through denying the claim that rational and moral reasons are commensurable. I used the rejection of the commensurability assumption to motivate my communitarian account of

⁷⁵ I should note that attempts to reduce morality to rationality, such as Gauthier's contractarian view in *Morals by Agreement*, are compatible with the communitarian moralist strategy because a successful reductionism will mean that agents' moral conduct can be entirely explained by their rational reasons to act. If a reductionist view turns out to be right, agents always have rational reasons to be moral. Such a view is compatible with my account because it does not situate purely moral reasons, derived from some overarching imperative to participate in moral projects regardless of one's particular projects, commitments, etc., at the level of agent conduct. Agents' reasons to do the moral thing will be their rational reasons to conduct themselves in particular ways. A successful reductionism will amount to an irrefutable argument for why agents always have good *rational* reasons to comply with morality.

morality. I argued that we make a mistake in thinking that moral reasons, like rational reasons, are reasons for agents to act. Instead I suggested that we see moral reasons as impersonal reasons for a community to adopt a certain kind of social organization that has good human intercourse as its goal, best realized by institutions that minimize interference in the personal projects of its members. On my account, agents' reasons to act are always rational reasons and the imperatives placed on agents to engage in moral conduct are entirely determined by the strength of their particular preferences to be moral.

According to my communitarian moralism, agents are not required to both qualitatively and quantitatively morally maximize their actions. Rather, as I argued in chapter eight, it is their community that is required to always both qualitatively and quantitatively strive to become better. The participation in community-bettering projects of agents comes entirely from the particular features of their lives and their particular character, which is bound up with the personal projects that they care about. I showed that the maximizing intuition is not about the agents' conduct both because it would have to ignore all of the particularizing features of agents' circumstances that give rise to the many different ways that agents choose to participate in moral projects and because it would require agents to implausibly prefer to participate in moral projects over non-moral ones.

In the final two chapters I addressed two important worries about my communitarian account. In chapter nine I responded to the charge that the impersonal stance's impartiality is threatened because it will always smuggle in some values of whoever adopts the stance, or from the community that she purports to represent, which

will inevitably bias her deliberations. I argued first that because the impersonal stance is adopted reflexively, if an agent adopts it on behalf of her community's values, the stance remains impartial toward the members of her community. I then suggested that when the values that she brings into the impersonal stance issue from her personal convictions, rather than the convictions of the community as a whole, there is opportunity for criticism and revision of the outcome of her deliberation.

In chapter ten I argued for how we can morally criticize agent conduct on a communitarian account of morality. I suggested that moral criticism can amount to a simple rational criticism of an agent's actions in some cases and an opportunity for more complex moral education in others. We saw that agents who hold inconsistent personal projects, some of which violate the community's moral institutions, can be rationally criticized for their immoral conduct. When their preference set is consistent but has immoral preferences, we can try to teach them to cultivate the right sorts of moral preferences. I ended by suggesting that those cases in which agents are unable to revise their preferences through rational criticism or moral education are in some sense lost causes. But while these cases are worrying, I think that they are quite rare, and that some other individualistic account of morality will do no better than communitarian moralism by trying to remedy the morally worrying preference sets through appeals to universal, overriding moral obligations. While some might charitably think that an individualistic account can give good grounds for agents to submit themselves to moral demands, such an account does worse than my communitarian moralism because it cannot explain why agents should *care* about being moral. In fact, the communitarian moralist's account of

the possibility of morally criticizing agents' actions is one of its strengths because, when such criticisms are successful, the agent will have a personal stake in acting morally.

BIBLIOGRAPHY

- Aristotle. "Nicomachean Ethics: Book IV." Accessed 08/08, 2013.
<http://classics.mit.edu/Aristotle/nicomachaen.4.iv.html>.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Brink, David. 1988. "Sidgwick's Dualism of Practical Reason." *Australasian Journal of Philosophy* 66 (1): 291-307.
- Darwall, Stephen. 2006. *The Second-Person Standpoint*. Cambridge: Harvard University Press.
- Cullity, Garrett. 2009. "Demandingness and Arguments from Presupposition." In *The Problem of Moral Demandingness: New Philosophical Essays*, edited by Timothy Chappell, 8-34. Basingstoke: Palgrave Macmillan.
- Gauthier, David. 1986. *Morals by Agreement*. Oxford: Oxford University Press.
- Hobbes, Thomas. *Leviathan*. Accessed 08/09, 2013.
<http://oregonstate.edu/instruct/phl302/texts/hobbes/leviathan-contents.html>.
- Hurley, Paul. 2009. *Beyond Consequentialism*. Oxford: Oxford University Press.
- Hurley, Paul. 2006. "Does Consequentialism Make Too Many Demands, Or None at All?" *Ethics* 116 (4): 680-706.
- Kant, Immanuel. 2005. *Groundwork for the Metaphysics of Morals*, edited by Jonathan Bennett. Accessed August 9, 2013.
<http://www.earlymoderntexts.com/pdf/kantgrou.pdf>.
- Kant, Immanuel. "Metaphysics of Morals." In *Practical Philosophy*. Edited by J. Gregor, Mary. Translated by J. Gregor. Mary, 353-604. Cambridge: Cambridge University Press.
- Kavka, Gregory. 1984. "A Reconciliation Project." In *Morality, Reason and Truth: New Essays on the Foundations of Ethics*, edited by D. Copp and D. Zimmerman, 101-113. Totowa: Rowman and Allanheld.
- MacIntosh, Duncan. 1991. "Preference's Progress: Rational Self-Alteration and the Rationality of Morality." *Dialogue: Canadian Philosophical Review* XXX: 3-32.
- Murphy, Liam. 1993. "The Demands of Beneficence." *Philosophy and Public Affairs* 22 (4): 267-92.
- Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford University Press.

- Orsi, Francesco. 2008. "The Dualism of Practical Reason: Some Interpretations and Responses." *Etica & Politica / Ethics & Politics* X (2): 19-41.
- Rawls, John. 1996. "A Theory of Justice." In *Reason at Work* 3rd edition, edited by Stephen M. Cahn. Orlando: Harcourt & Brace Company.
- Scanlon, T. M. 1998. *What we Owe to each Other*. Cambridge: The Belknap Press of Harvard University Press.
- Scheffler, Samuel. 1992. "Prerogatives without Restrictions." *Philosophical Perspectives* 6 (Ethics): 377-397.
- Scherkoske, Greg. 2013. "Whither Integrity II: Integrity and Impartial Morality." *Philosophy Compass* 8 (1): 40-52.
- Sidgwick, Henry. 2011. *The Methods of Ethics*, edited by Jonathan Bennett. Accessed April 25, 2013. <http://www.earlymoderntexts.com/pdf/sidgmeth.pdf>.
- Stockdale, Katie. 2013. "Collective Resentment." *Social Theory and Practice* 39 (3): 501-521.
- Tessman, Lisa. 2005. *Burdened Virtues: Virtue Ethics for Liberatory Struggles*. Oxford: Oxford University Press.
- Walker, Margaret Urban. 2003. *Moral Context*. Lanham: Rowman & Littlefield Publishers.
- Williams, Bernard. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*, edited by J. J. C. Smart and Bernard Williams, 75-150. Cambridge: Cambridge University Press.
- Wolf, Susan. 1982. "Moral Saints." *The Journal of Philosophy* LXXIX (8): 419-439.