

Cost and the Limits of the Obligation to Aid

By Josh Gardner
University of Michigan Honors Thesis, Spring 2010

Advisor: Sarah Buss
Second Reader: Daniel Jacobson

Acknowledgements

First, I would like to thank my parents, Richard and JoAnn Gardner, for their continued and enthusiastic support. Had it not been for your encouragement, I may have never discovered my interest in philosophy in the first place, let alone completed this thesis.

I also want to sincerely thank my advisor, Sarah Buss. Thank you for teaching me how to read, write, and think carefully – and thank you for reading and thinking carefully about the countless outlines, proposals, questions, and drafts that peppered your inbox throughout the year. Thank you for asking difficult questions. Thank you also for asking simple questions—sometimes these were the most difficult and the most rewarding to answer. Thank you for being honest with me and for holding me to a high standard. Working with you made this challenging project also an extremely enjoyable one. I will continue to apply the lessons I have learned while working with you in my constant encounters with problems requiring careful, disciplined thought.

I would also like to thank Victor Caston for his keen and thoughtful feedback, encouragement, and patience while I developed and refined the project. Had it not been for your valuable advice about identifying a concrete problem and, after I thought I had, how to approach solving an issue that at first seemed too overwhelming to approach, my project might still be a collection of scattered thoughts and notes.

I would also like to thank the other Class of 2010 Philosophy Honors Concentrators for their thoughtful discussion.

I would like to thank the entire philosophy department at the University of Michigan. The inquisitive environment that the department fosters, one that is compelled by interesting and difficult philosophical problems, gave me the skills I needed to complete this project. The department's collective passion for thinking hard is contagious. My classes in the department provided a strong foundation for philosophical inquiry, and the instruction and feedback I received from members of the department was invaluable in completing my thesis. In particular, I wish to thank Daniel Jacobson for his thoughtful and challenging feedback and discussion of the final essay.

Finally, I would like to thank Molly Mahony at the Tanner Library for her assistance throughout the year. Most of the books and articles I read for this project were found using the research tools that Molly pointed out to me. Her helpful pointers truly improved the work and made the research process much more accessible and thorough.

I. Introduction: Cost and the Demands of Morality

It seems clear that we have a moral obligation to aid others. But, given that helping others is often just one among our many ends, then we also accept that morality and our self-interest will, in some cases, not be in harmony. Morality's demands will sometimes require immense sacrifices of our own resources and projects in order to aid others in great need—and for those of us who have not only the knowledge but also the means to do something to relieve suffering around the globe, such demands will occur frequently. It is difficult to determine, however, whether the imposition of costly demands is ever a *problem* for morality, or whether it is simply a part of being moral that we must learn to cope with, however difficult. In this essay, I will explore the possibility that there is a limit to what we can be morally required to sacrifice in order to aid others. Judith Lichtenberg observes that “once we admit duties to aid into the moral realm they threaten to take over and invade our lives: it is hard to draw a line that will prevent them from becoming relentlessly demanding”.¹ I will begin to sketch this line.

There is a great deal of intuitive support for some limit to moral requirement. To many, the claim that we can be morally required to make great sacrifices in order to aid others is “absurd, and...surely wrong”.² Indeed, I will argue that ordinary morality—the collective name given to the set of our everyday moral judgments—seems committed to defending the idea that there is a limit to what we can be morally required to sacrifice. This limit, however, is often vaguely understood, and determining exactly what justifies it will be a primary concern of this essay. I will argue that the question of whether there is a limit to our obligation to aid others is, ultimately, a question about whether *other*

¹ “Negative Duties, Positive Duties, and the “New Harms”,” 1.

² Joseph Raz, “A Morality Fit for Humans”, 1297.

sources of moral value limit this obligation. The search for a limit to moral requirement strictly in terms of the demandingness of morality—the sacrifices it imposes—is futile without reference to something important that is being given up, because morality’s authority to demand sacrifices of us is potentially unlimited; there is no *general* limit to what it can demand of us. I will argue, however, that the moral significance of our agency—our capacity to act for reasons, to form and pursue plans and attachments over time—raises special concerns which will, indeed, limit what morality can demand of us in order to aid others, in many cases. Certain sources of moral value, such as the projects and relationships that shape our lives and give them meaning for us, can only be recognized in ways that limit our moral requirements to aid others. While I do not attempt to give an exhaustive account of agency, I believe that even the basic account I offer is supported by the moral judgments we wish to make, as well as our notions of the source of our moral requirements. Even with such a limit, however, it is likely that morality will still require sacrifices of us—and sometimes very great ones—although this will happen significantly less than under “extremist” moral systems.

The proposal that there is a limit to our moral requirement to aid others has met with a great deal of resistance, however. In particular, I will discuss the arguments of Shelly Kagan, who insists that the only acceptable moral system is an “extremist” one under which agents’ duty to promote the impartial good is their only moral obligation. Under Kagan’s system, agents have an unlimited obligation to promote the impartial good—morality can and will often require immense sacrifices of agents’ resources and interests. Kagan holds that any moral system which fails to maximize agents’ contribution to the impartially-conceived good is fundamentally inadequate, and he

presents some formidable challenges to any attempt to justify a limit on moral requirement. In particular, I will discuss Kagan's claim that any limit on moral requirement must point to some source of moral value *not* recognized from the impartial point of view. Kagan believes that this is an insurmountable challenge, but I will argue that our agency is indeed such a source of value—it cannot be given its full and necessary recognition from the impartial point of view.

I will begin my discussion, in the next section, by outlining some basic features of moral judgment that will prove essential in understanding how a limit on moral requirement might function, and what exactly the rejection of such a limit amounts to. I then consider some proposals for justifying such a limit before offering my own.

Objections to highly demanding moral theories fall into two basic categories: those which object to morality's demandingness itself, holding that extreme demands are unacceptable for some reason; and those which see demandingness as a mere *symptom* of a deeper problem with such theories. Ultimately, I will argue that approaches of the first type fail: a moral theory which considers all morally relevant factors cannot be “too demanding;” to say otherwise is to misunderstand morality itself. I will argue for a version of the second approach: excessive demandingness might indicate a theory's failure to account for important sources of moral value—a failure that can, however, be corrected within the existing framework of impartial moral systems, but requires the rejection of Kagan's extremism. Specifically, I will argue that the only plausible variant of the “appeal to cost”—a limit on our obligations to aid others—is supported by a parallel understanding of the foundation of these moral obligations: both emphasize the moral

significance of *agency*, which morality seeks to promote as a source of value and which its demands cannot, therefore, encroach upon.

II. Understanding Moral Deliberation

Much of the difficulty of determining whether there exists a legitimate basis for any cost-based limitation on moral requirement concerns the prior difficulty of identifying how we make moral judgments, in at least some basic sense. Here I highlight some basic elements of how we determine what we are morally required to do which are fundamental to understanding the thought of a limit on moral requirement. These elements are: the responsiveness of morality to moral reasons, the pro tanto reason to promote the good, constraints and options, and the appeal to cost. I will refer back to this basic framework of moral thinking throughout my analysis, and these concepts aid greatly in my diagnosis of the problem with highly demanding moral theories and my recommendation about how to repair such theories. It should be noted that while this analysis draws heavily from philosophical discussions of “ordinary morality,” the basic process outlined here is applicable to most, if not all, moral theories.³

The first, and perhaps most important, feature of moral deliberation is its recognition of the morally good action as that which is supported by the balance of morally relevant reasons. In this essay, I adopt a primitive interpretation of a reason, taken from Scanlon, as “a consideration that counts in favor of” an action.⁴ I use Scanlon’s broad interpretation in order to open up the category of reasons to include a wide variety of considerations that may be excluded by more restrictive interpretations but which seem to

³ This account of ordinary morality relies especially, but not entirely, upon features highlighted by Kagan (1989).

⁴ See T.M. Scanlon (1998), “What We Owe to Each Other” for a discussion of this type of reason.

be significant for moral discussions. Furthermore, Scanlon's interpretation recognizes that there may be a variety of competing reasons for and against any given action, and that these reasons are not exclusively generated by our own ends or desires. A morally relevant reason, on this account, would thus simply be a *moral* consideration which counts in favor of doing something—one which is attached to a source of distinctly moral value.⁵ As I will discuss in detail later, I believe that highly demanding moral theories often issue such costly requirements as a result of their failure to recognize certain reasons as morally relevant. My primary point here, though, is simply that the weighing of moral reasons forms the basis of moral judgments. If we act contrary to the balance of moral reasons, then, we are acting immorally; and if we fail to recognize all the necessary reasons, then, being unaware of the true balance of reasons, we are likely to act contrary to it.

This recognizes our understanding of morality as a project closely linked to rationality, as both require responsiveness to reasons. My emphasis on moral reasons of course serves to emphasize that there exist different *kinds* of reasons, not all of which are relevant to moral deliberation. A police officer may enjoy the feeling of *schadenfreude* she experiences when she pulls a car over even though the driver did not violate any traffic laws, and may therefore have a *reason* to do so, but this reason is not relevant to her moral duty to fairly enforce the law: not all reasons are morally relevant reasons, and an important part of moral judgment is deciding which reasons are morally relevant. An agent can have good practical or self-interested reasons to perform a given action without

⁵ Deciding what is an explicitly *moral* consideration will emerge as an important issue, and I will argue that most impartial theories have failed to take account of the full spectrum of moral reasons. I will not, however, offer a direct analysis of what distinguishes moral reasons from other kinds of reasons.

having any moral reason to do so; she might even have moral reason(s) to perform an action when stronger moral reason(s) oppose her performing that action.

Another important feature of moral judgments is that they generate a specific type of reason, often called a “pro tanto” reason, in favor of actions which promote the good (however that good may be conceived).⁶ While we may not always (consciously) acknowledge this reason, it is important that we recognize it here: the existence of the pro tanto reason reveals that the moral status of our available actions is always relevant to our decisions about what to do. While the pro tanto reason may not always be *decisive*—it will, of course, contend with other kinds of reasons for action—and while it is not overriding on most conceptions of morality, recognizing its existence helps to explain ordinary moral reasoning and, as we will see later, explains why we would also object to moral theories that fail to be demanding enough. It is worth noting that morality would seem empty without it: a moral system which generated judgments, but whose judgments provided no reasons for acting, would be odd, useless, and quite unlike anything we think of as “morality.”

Additionally, two complimentary features of moral deliberation are essential in understanding how limitations on moral requirement fit into the process of moral judgment. *Constraints* and *options* are both important parts of all but the most demanding conceptions of morality, and in this essay I will use the term “extremism” to refer to moral theories that recognize the two prior features (a moral system based on the balance of reasons and the presence of the pro tanto reason), but do not recognize the existence of

⁶ The term ‘pro tanto’ has been used in a variety of distinct contexts, and its use here should not be confused with other uses. Again, I borrow this concept from Kagan(1989), but it is used widely throughout the relevant literature. Here I intentionally avoid making any claims about what constitutes the overall good in the hope of making this account of moral deliberation applicable to a wide variety of moral theories.

constraints and options. *Constraints* are limitations on what actions an agent can be morally permitted to do: for example, a commonly accepted constraint on murder holds that no human is ever morally permitted to kill another innocent human. A constraint is thus an *exclusionary* reason: constraints categorically exclude certain types of actions, no matter what reasons she may have for doing so—the actions are “off the table” as far as we are concerned. Kagan rejects all constraints because they seem to prohibit an agent from performing a given act (such as killing) even in order to prevent multiple acts of the same type—that is, they seem to morally condemn actions which would, according to the consequentialist, actually maximize the overall good. However, I will later argue that, given the moral significance of agency, Kagan is wrong: in order to recognize the moral significance of agency, impartial theories must adopt an agency-based constraint.

Options allow agents to choose to pursue their own interests within the limitations of constraints. When (and if) we have options, we are free to perform any action not ruled out by constraints. Options, then, give agents permission to act contrary to the pro tanto reason to promote the good, allowing them to act in ways which do not maximize the overall good. Of course, agents with options may still choose to maximize the good—provided doing so does not violate any constraints—but they are not morally required to do so; such action would be supererogatory. On a consequentialist conception, options are simply permissions to pursue actions which lead to outcomes that are less than optimal on an objective ranking of available outcomes. It is easy to see why constraints and options might be required to complete the picture for most moral theories: without them, and with only the beliefs that the morally-required act is that which is supported by the balance of morally decisive reasons and that there is a pro tanto reason to perform the

goodness-maximizing act, agents would always be morally required to maximize the overall good. Such a system would be an *extremist* one, and hardly resembles commonly held beliefs about our moral obligations. Constraints and options place limitations on what we can be morally required and permitted to do, and ordinary morality seems committed to their existence.

A final, more specific element to which ordinary morality is committed is “the appeal to cost.” Given the content of our ordinary moral judgments, it seems that the ordinary moralist is committed to endorsing a principle which holds that the cost of performing a given action is relevant to whether an agent can be morally required to perform it. I will attempt to identify this principle in greater detail in the following sections, but wish to note here that ordinary morality seems committed to the existence of such a principle: not only would it grant a wide variety of options by eliminating moral requirements to perform actions that are excessively costly to perform, but it seems that such a principle would yield demands that match up with ordinary morality’s judgments about *which* cases we have options in. The appeal to cost does not grant agents unlimited license to pursue their own interests and to fully neglect the demands of morality; instead, the appeal rules out particularly costly requirements, leaving those obligations that fall within some acceptable range of demandingness. Some obligations certainly still exist in the world of the appeal, as only the particularly costly demands are excluded. This seems to match up with the judgments of ordinary morality: far from seeking to deny that any moral obligation exists, the moderate merely argues that extremely costly moral demands are not acceptable. While considerations of costliness may not always factor explicitly or consciously in ethical judgment, in most cases—especially when the demands made of us

are especially great—we do take cost into consideration, and it is often cited as a grounds for rejecting particularly burdensome moral requirements. While there may be a variety of other considerations that can limit our moral requirement, ordinary morality seems fully committed to the claim that cost is one of them—and one of the most important. Furthermore, it is of note that we seem to truly *believe* that cost is of moral relevance: persons highlighting considerations of cost usually do not appear to be merely rationalizing, but are convinced that they are pointing to something of genuine moral significance.

III. The Appeal to Cost

In this section I will offer an initial description of the appeal to cost. I will consider possible justifications for the appeal in the following section. In its most general formulation, the appeal to cost amounts to the claim that the costs imposed on individual agents by morality's demands are themselves morally relevant, and that any moral theory which is not sensitive to these limits must therefore be rejected. I take the appeal to cost to be one of the defining features of ordinary morality. Indeed, even if agents do not frequently point to the costly demands of certain moral systems as their reason for rejecting them, a great deal of our actions (and the corresponding moral judgments that permit or justify these actions) seem to reflect this belief that there is a *limit* (one low enough that it excludes many requirements that would be generated by an extremist moral system) to what morality can demand of an agent.⁷ The appeal to cost adds *options*

⁷ I should specify that in this essay I am primarily interested in considerations of cost primarily in cases of (positive and negative) aid to others, but my analysis applies to moral requirement more broadly. I do not

to extremist moral conceptions⁸, giving agents permissions not to maximize their contribution to the overall good when doing so would be unacceptably costly. The appeal to cost does not, of course, hold that it would be morally wrong for agents to undertake great sacrifices to promote the overall good; it merely holds that they cannot be morally required to do so. The appeal to cost demonstrates how the balance of reasons considered in moral judgments can be significantly altered when different reasons are recognized: when we add to extremism the reasons generated by the appeal to cost (thus producing ordinary morality), many of our moral duties to provide aid to others are eliminated, and those still present are substantially less costly.

Beyond recognizing that the appeal to cost eliminates moral requirements which exceed a certain morally relevant level of costliness, identifying exactly what this “morally-relevant” limit is, and how exactly it is to be formulated, is a challenging and complicated task, and identifying and testing more specific formulations of the appeal will occupy much of this work. Even after carefully laying out the possible formulations of the appeal to cost, however, the moral moderate is still faced with the possibility that she has attempted to give form to an ultimately formless intuition, seeking a standard where none exists. Here I will begin by sketching the basic features of any version of the appeal, and will point out the implications for any moral system of granting such an

take it that morality *only* requires us to aid others—indeed, I think there is a strong basis for a moral requirement to develop ourselves. I do not discuss this duty much here, though, because this is hardly a controversial duty—most persons are only too happy to recognize a moral duty of self-development. I think this duty equally interesting as cases of aid, but my discussion here is confined to the latter. Both, though, have a common foundation, as I argue later in the essay.

⁸ While, as mentioned above, I use the term “extremist” here to broadly indicate any moral theory that always requires agents to maximize the overall good, these moral theories might also be characterized as adopting a “maximizing” approach; Portmore (2007) discusses this maximizing approach as the feature that makes act-utilitarianism so compelling, even for those who reject it. I do not use the term “extremism” to denote the particular form of consequentialism endorsed by Kagan under the same name, but my broader use of the term does cover a theory like Kagan’s. For a discussion of Kagan’s extremism, see *The Limits of Morality*.

appeal. After uncovering some necessary features, benefits, and implications of the appeal, I will consider how the defender of ordinary morality (hereafter the Defender) might justify it.

I take it that the costs highlighted by the appeal are best understood in terms of the burdens moral demands place on agents' ability to pursue their own interests. If it is to justify ordinary morality, the appeal must be concerned with agent's interests *as a whole*, and not only their *self*-interest: for example, we care about the well-being of our friends and family, even when this is not in our self-interest. Furthermore, when morality's demands merely drain agents' resources, not directly affecting their well-being nor that of their loved ones, such demands are still considered objectionable—thus, this draining of resources itself must somehow be morally relevant. This can also be understood in terms of interests: if not spent complying with morality's demands, resources such as time, energy, and money would otherwise be used to promote agents' own interests. The loss of these resources is thus relevant insofar as it prevents agents from using these resources to promote their own interests. Because it relies on an agent-relative conception of cost—one based on a demand's cost for a particular agent, given her interests—this interpretation of cost has the consequence that a potential moral demand, such as “donate \$500 to local literacy programs for the poor,” will be more or less costly for particular agents depending on their interests. This demand might mean that, as a passionate runner of marathons, I would have to sacrifice hours that I wished to spend running in order to earn the \$500, making it very demanding for me. For someone whose greatest passion was literacy promotion, however, this demand could cost nothing at all. So more generally, for agents whose only priority is to promote the overall good, even an extremist moral

system might not actually impose “costs” in this sense. Such costs only emerge when the demands of impartial morality and agents’ own interests diverge. While the appeal to cost locates the problematic demands on the individual level, it requires that we take steps—such as introducing options and constraints—to modify our moral theory itself to prevent such demands.

Understanding exactly what is meant by “demand” and “cost” in this context is essential to understanding objections to moral theories which impose either in excess. Here I restrict the term “demands” to refer only to moral requirements that impose *costs* on agents. Moral theories which make a litany of commands but which do not impose any great loss of resources—such as a stringent form of ordinary morality which *requires* that well-off nations do something, but only very little, for those in great need, as they do now—are not, in this sense, demanding (although we might certainly identify other problems with such a theory). Instead, only those theories whose requirements have significant negative effects on agents’ ability to pursue their own interests by draining their resources—for instance, one which requires me to dedicate my valued time spent training for a marathon to an activity which better promotes social welfare—are subject to objections for being too demanding. According to the appeal to cost so understood, then, it is only the *losses* imposed on agents, not the actual extent or stringency of a moral system’s mandates, which are objectionable.

We can also see that the appeal to cost is concerned with *relative* losses for any given agent, and not on their falling below any absolute level of well-being as a result of moral demands. We think that the well-off and the not so well-off person would both be justified in objecting to a requirement to donate most of their disposable income to

provide emergency relief, even though the well-off person would still have an absolute level of well-being much greater than the less well-off person did before donating anything to charity. The only way both agents could be understood as making the same, coherent objection would be if they objected to the relative losses imposed on them, not the absolute level of well-being they were reduced to by the costs of doing their moral duty.⁹

Additionally, a note about the applicability of the appeal to various types of moral theories. Discussions of the appeal to cost have typically centered upon its relation to consequentialism, especially “common-sense consequentialism.” The appeal is discussed as one of the primary features distinguishing “ordinary morality” from extremist consequentialism. I would like to point out, though, that the appeal’s role of placing a limit on moral requirement could just as well be applied to nonconsequentialist moral theories. As Murphy has noted, the appeal to cost can be raised in objection to any theory that includes a principle requiring agents to promote the good.¹⁰ This means that we cannot simply reject consequentialism, citing its overdemandingness, and hope to find a theory that might be able to better address these concerns. No matter where we turn to identify our moral duties to aid others, if we have any at all, then the demandingness of moral theories must be dealt with. Rejecting consequentialism will not suffice to escape

⁹ For a discussion of the distinction between losses and absolute level in the problem of overdemandingness, see Murphy, *Moral Demands in Nonideal Theory*, 20-21.

¹⁰ See Murphy (1993): “Far from being a problem just for theories that include or are constituted by the Simple Principle, we have a problem for all moral theories that include a principle requiring us to promote the good. Some writers concerned about over-demandingness simply point out that consequentialism, because it does not limit its demands, is unacceptable. But while this may be true, it misleadingly suggests that other moral conceptions are better off in this respect; in the absence of any account of how to understand and justify a limit to the demands of a principle of beneficence, they are not” (273). Furthermore, if the project of “consequentializing” moral theories succeeds, then it seems that the appeal to cost is one that can be made relevant to any moral theory in this way as well: one need only insist that the cost to agents of complying with the theory’s demands is relevant to the deontic ordering of outcomes on that theory. See Portmore, Douglas W., “Consequentializing Moral Theories” (2007).

the problem of explaining what, if anything, is wrong with moral theories that demand much of agents, and how we can modify our theories to accommodate this concern.

A final note about a feature of the appeal that I will not seek to elucidate: to be of any practical use, any formulation of the appeal to cost would have to contain some relatively precise principle identifying where the morally-relevant threshold for cost lies.¹¹ In this essay, I am primarily only interested in determining *whether* such a limit on moral requirement exists. It is possible that, even if the appeal to cost or some similar limit on moral obligation is justified, we cannot identify its limits with any specificity—but if such a limit does exist, I leave the challenge of identifying its precise boundary to others.

IV. Defending the Appeal

I now turn to the difficult task of inventorying possible justifications for the appeal to cost. This assumes, of course, that the appeal does indeed require some defense, that we must give reasons for accepting this claim. As with many moral principles, however, there remains the possibility that it is somehow intrinsically or trivially true—I will briefly consider this as a strategy of justification—but my assumption here is that, given the intimate connection between morality and rationality, our moral principles rest on reasons which can be identified and tested. Ultimately, I will argue that there exists no legitimate foundation for the appeal to cost itself, but that excessively costly demands

¹¹ As Judith Lichtenberg (2010) points out, The “Hand Rule” formulated by Judge Learned Hand in *U.S. v. Carroll Towing Co.* is a rather plausible example of such a precise limit; its continued use in American jurisprudence also suggests that it might be a useful principle for identifying the limits to moral requirements, at least in a certain type of case (i.e., in cases where persons are required to avoid immoral negligence of others).

might be a *symptom* of another problem with such theories—their failure to recognize important moral reasons related to agency.

When considering justifications for the appeal, it is also of use to identify what cannot function as a justification. Here I am specifically interested in making clear that the mere fact that agents do indeed resort to the appeal to cost, and do so with surprising frequency, does not justify its use. Facts and fallacies are just as often repeated and find their way into our habits. We must look beyond the fact that we do indeed use the appeal to discover what basis, if any, it might have.

I begin the search for a justification by considering the most straightforward approach that the Defender could take to justify the appeal: that it is just a fact that morality cannot demand too much from us, and that this is a self-evident claim that we need not justify. This claim that there is a basic, fundamental limit to what morality can demand of an agent—let us label it the “straightforward justification”—is indeed how many of us might be inclined to respond when questioned about why we oppose highly demanding moral theories. Much like an obviously true mathematical principle, the Defender might claim that we simply “know” that there is a limit to how much morality can require us to give up in order to perform our duty. Murphy (1993) takes seriously this claim that there is no underlying justification of objections to extremely demanding moral theories. Murphy specifically rejects Kagan's insistence on the “need” for some deeper justification for this objection, claiming that the appeal to cost has and requires no “underlying rationale” and can be taken at face value—as the claim that there is a limit to what morality can demand of us in order to promote the good.¹²

¹² Murphy, *The Demands of Beneficence*, 273-274. Murphy calls this the “demandingness objection.” Scheffler (1986) also discusses this view.

If all the appeal to cost rests on is the straightforward justification, however, then the Defender is in trouble. Murphy rejects the straightforward justification, pointing to the difficulty of finding any specific formulation of the argument which is usable in a consequentialist calculation and which maintains its simple, obvious appeal. There are other reasons to reject the straightforward justification as well: such claims cannot lack justification “all the way down”; somewhere, there must be an underlying argument. Otherwise, the straightforward justification seems to require that we abandon the connection between morality and rationality—an unacceptable consequence.

The straightforward justification might rely on the assumption that our intuitions can give us important insights into real moral laws or facts. Kagan, however, specifically rejects this intuitionistic approach to ethics, and rightly so.¹³ Relying on our intuitions alone is especially problematic in the case of the appeal to cost, because we have conflicting intuitions: we not only believe that there is a limit to moral requirement, but we also believe that morality requires us to promote the good—which may require great sacrifice. While intuitive support certainly seems to count as a consideration in favor of (a *reason* for) adopting a principle, and while we must rely on our intuitions at some point, in cases where our intuitions conflict they do not help resolve the issue. Because it is unclear how this intuition might (or might not) resist our other intuitions about the potential demandingness of moral requirements, then, and because we have already recognized that morality is not a project grounded in self-interest, the burden shifts to the Defender, who must explain why cost is relevant to moral requirement. The straightforward justification alone is thus insufficient; the Defender must look elsewhere to justify the appeal to cost.

¹³ Kagan, *The Limits of Morality*. See “Moral Methodology”, 11-15.

Murphy himself rejects the straightforward justification in favor of his own argument, which supports a slightly different limit on moral requirement. Murphy claims that we don't actually have any independent basis for objecting to theories simply in virtue of their demandingness—certain costly demands, he thinks, are acceptable. He holds that what we ought to object to is theories whose demands of agents increase as others fail to perform their moral duties, requiring agents to pick up the slack for their fellow agents. He believes that this objection “expresses what was worrying us about [theories that require agents to maximize the good] all along. Thus we could argue that the over-demandingness objection [i.e., the appeal to cost] does not need to be solved on its own terms, but rather dissolved—we never were that worried about the *amount* of the demands; instead, we were all along concerned about doing the work of others”.¹⁴

Murphy emphasizes that morality is a “cooperative” project which requires the participation of everyone involved. Because morality is such a project, he argues, demands on agents cannot increase as other agents' compliance with morality decreases: when others fail to cooperate, we cannot be morally required to carry out their moral duties in addition to our own. Murphy's arguments might be seen as a more sophisticated appeal to cost—one which actually locates a specific undesirable type of costliness and gives some justification for why it is objectionable. According to Murphy, what we object to in extremist theories is the fact that their demands are inflated by the world of “partial compliance” we occupy—all agents do not fully comply with their moral obligations. Extremist moral theories, then, unacceptably force those agents that do comply with morality to endure significantly greater demands as a result. In a world of full compliance, extremism's demands would not, Murphy believes, be objectionable—

¹⁴ Murphy, “The Demands of Beneficence”. 289. See also *Moral Demands in Nonideal Theory*.

even if they were quite costly—but given the reality of only partial compliance, extremist moral systems make unfair demands of those agents that do attempt to comply. So, if morality’s demands were adjusted such that our own obligations did not depend on how much others did, then according to Murphy, they would then be unproblematic no matter how costly they were. By recognizing that the underlying problem with excessive demandingness is actually a problem with moral systems that fail to account for the effects of compliance on moral demands, Murphy “dissolves” the problem of overdemandingness. To resolve this problem, he suggests a “compliance condition,” which requires that our obligations not be affected by others’ performance of their duties—that we each only be required to do what our moral duty would be in a world of *full* compliance.

I agree with Murphy that we need to “dissolve” the problem of demandingness—to focus on *specific problems* with the demands that explain why the costs are morally relevant. As it stands, the appeal to cost does seem insufficiently general, and certain types of costly requirements might be perfectly acceptable, while others might be truly objectionable—and when we dissolve it, we can clarify which considerations might justify some limit to moral requirement. However, I do not think that Murphy’s attempts to dissolve this problem identify the most important way in which such demands are problematic. While we may indeed want our moral systems to be sensitive to the effects of compliance, the demands of such a system would still be objectionably costly, albeit in other ways. If this is so, then the demandingness objection isn’t fully, or even *mostly*, resolved by Murphy’s proposal.

Under Murphy's proposal, moral theories might be significantly less demanding than unconstrained extremism, because agents' obligations would not be determined purely by the actual state of the world (which is one where many agents fail to meet their moral obligations), but instead by how much agents would need to do if *everyone* were meeting their moral obligations. But, as Murphy recognizes, his theory would still probably be extremely demanding: there exists a great deal of widespread, preventable suffering in the world, and for many individuals, morality would still make costly demands, even if everyone else did their moral duty.

It seems that agents that are still required to sacrifice a great deal, even under Murphy's compliance condition, would raise the same, highly plausible objection to Murphy's theory that they raise in response to extremist moral systems: it still makes very costly demands of agents; demands which appear no less objectionable merely because they are really our *own* obligations instead of those we inherited because of others' moral shortcomings. If the devoted but financially well-off artist objected to extremism on the grounds that it significantly limited his ability to pursue his greatest interests, why would he object any less to Murphy's account? Murphy's account still imposes costly demands, and simply pointing out that these demands are only "our fair share" does not appear to make the costly requirements any less problematic. The fact that Murphy's account demands nearly as much, but merely for a different reason, still meets with much of the intuitive force of the appeal to cost.

This is not to deny that Murphy has identified a problem with extremist moral systems; I instead believe that there are much larger problems with extremism that Murphy's account fails to address. It seems that, as Murphy mentions is a possibility,

even if his account solves a very real problem with extremism—its insensitivity to compliance—his account still fails to resolve their problem with overdemandingness. So, Murphy’s attempts might actually solve a *different* problem with moral demands, while leaving the problem highlighted by the appeal to cost unsolved (or undissolved, as the case may be). Murphy’s account fails to explain why, even when the effects of compliance are mitigated, his theory might still be problematic, and he does little to prove that there aren’t other problems with highly demanding moral theories besides their insensitivity to compliance. It seems quite likely that there is more than one problem with extremist moral theories, and Murphy’s apparent resolution of problems of compliance misses what I will soon argue is a more important problem with highly demanding moral theories—their demands reflect a failure to properly account for the moral significance of agency.

Another consideration that the Defender might point to in order to justify the appeal to cost is that highly demanding moral theories *alienate* us from our own aims and projects, and that this alienation is unacceptable. Such a strategy seems quite plausible: the impartiality of extremist moral systems does give strikingly little regard to questions of integrity and personal identity, and this could be what lies at the root of objections to their demands. Given the separation of morality and self-interest, morality’s demands will sometimes alienate us from our own interests, commitments, and values in order to act morally (provided our ends aren’t in perfect harmony with morality). Morality will sometimes require us to abandon our own desires to take better care of our families, to spend time with friends, or to dedicate our income to supporting valued pursuits like music or marathoning which extremist accounts would prohibit whenever a morally

superior action is available. In a globalized world that provides us with both the knowledge and the means to relieve the great suffering of others (and which also often implicates us in their suffering in increasingly complex ways), such dissonance between our own projects and morality's demands may occur almost all of the time—resulting in frequent alienation as a result of morality's demands.

However, while the alienation approach identifies an important concern about moral demands, this concern alone is not sufficient to justify a limit on moral requirements to aid others—but it is correct to highlight the moral significance of agency, and of the integrity that is central to the expression of one's agency. Despite the initial plausibility of giving an account of the appeal to cost in terms of alienation, this is not a feature of moral demandingness to which the Defender can universally object: as Railton points out, alienation is not always a negative effect of moral requirements.¹⁵ In many instances, alienation is a part of a process we find morally good and necessary for moral development: it is not inherently wrong that agents are often morally required to hold up their beliefs for examination, and even to abandon deeply-held values, despite the fact that such processes are indeed alienating. We think it good that slave-owners were alienated from their convictions about the moral permissiveness of their practices, and the fact that they were alienated from their belief that slaveholding was acceptable hardly carries any moral weight at all. Learning to accept when our commitments conflict with morality is an important, if difficult, part of being moral. While morality certainly ought to accommodate considerations of integrity and agency, which I will discuss later, it cannot be expected to honor them unconditionally—and the slave-owner example suggests that we would not want it to.

¹⁵ See “Alienation, Consequentialism, and the Demands of Morality.”

As Railton also notes, alienation is also not a universally negative consequence of moral demands because in certain circumstances, alienation is necessary to maintain our *autonomy*: unless we are willing to surrender even our most closely-held beliefs and values, or at least to thoughtfully evaluate and reconsider them, we fail to be autonomous and are controlled by these convictions. Alienation alone, then, is insufficient to justify any opposition to moral demands—even if, in many cases, the alienation that results from moral demands is of moral significance, sometimes it will be necessary to preserve our autonomy, which is also undoubtedly of moral significance.

The problem with the alienation objection, then, is that it fails to differentiate between different cases of alienation—those in which the alienation is somehow acceptable, and those in which it is not. What this seems to show is that the objection cannot object to these demands strictly insofar as they are alienating: even though alienation might, in many cases, indicate some problem with morality's demands, the alienation itself is not the problem. While alienation does seem to be an important consideration for moral demands to recognize, the objection as it stands fails to provide any real guidance on *how* these demands ought to recognize it. If we understand the objection as holding that *all* instances of alienating demands are unacceptable, then the objection fails, and it is unable to articulate any threshold for when these demands ought to be accepted despite the alienation they will cause. I will argue that this alienation can be understood in terms of our *agency*, and that an agency-based constraint on moral requirement will allow us to differentiate between cases in which alienation is acceptable and cases in which the alienation genuinely ought to limit morality's demands.

There is another way the Defender might point to considerations of moral autonomy to justify the appeal to cost. The Defender might argue that highly demanding moral systems remove agents' ability to choose from a wide variety of morally-permissible actions to such an extent that this actually constitutes a morally relevant decrease in their autonomy. I will call this the "confinement objection."¹⁶ Extremist moral theories certainly have the potential to be extremely confining: when agents are required to choose only those actions which maximize the overall good, they will often have few, if any, choices about how to act, and many of the available actions will be ones which they would not otherwise have chosen. The question, then, is whether this narrowing of options constitutes a *moral* problem for such theories. While Murphy does not explicitly characterize the confinement objection in terms of *moral* autonomy, this seems to be the only way such an objection could affect agents' moral requirements: the moral significance of the autonomy under attack allows it to oppose claims of moral duty; it generates *moral* reasons in favor of preserving it. Shiffrin (1991) and Slote (1985) have challenged highly demanding moral theories on the grounds of moral autonomy. As Shiffrin notes, this approach derives its justification for options "not from the subjective importance of an individual's projects, but by instead appealing to the value of the agent's having a range of morally permissible choices between projects and actions available to her".¹⁷ This highlights an important feature of this approach: it holds that agents are morally entitled to having an acceptable variety of morally permissible options regardless of what these actions are and even of their importance to the agent herself. The confinement objection insists that all moral demands must reflect that all agents are

¹⁶ The term "confinement objection" comes from Murphy, *Moral Demands in Nonideal Theory*, 26-33.

¹⁷ Shiffrin, "Moral Autonomy and Agent-Centred Options", 244-5.

always morally entitled to some degree of choice about which actions to perform, and that morality must *approve* of these actions.

Despite the plausibility of this approach, derived from its emphasis on the significance of our ability to mold our own lives in ways that are morally permissible, the confinement objection cannot alone justify an appeal to cost that will limit our moral obligations to aid others. It seems wrong to hold, as the confinement objection does, that it is always morally better for an agent to have a wide variety of morally permissible options than for her to have her choice highly restricted by morality: confinement is not always a bad thing; or at least it would sometimes be much worse to avoid confinement than to permit certain actions strictly to avoid confining the agent's choices. In fact, sometimes we think it *better* that certain options are simply not on the table. We do not think, for example, that it is problematic for a moral system to give the slaveholder only one permissible option—*not having slaves*—instead of allowing him to choose between two. If the confinement objection were correct, then there would be a serious problem with all moral constraints, as they are, by definition, confining. While it may be better for me to have many morally-acceptable options in any given situation than to have only a few, when this requires extending the range of permitted actions to include otherwise-wrong actions, there is an additional consequence—that of otherwise-wrong acts being presented to the agent (and possibly also performed) as if they were morally acceptable. If we alter morality purely for the sake of giving agents choice, thereby turning otherwise morally wrong acts into morally “permitted” ones so the agent is offered a choice, then it seems that the result—an agent given options to perform otherwise morally wrong acts for the sake of her choosing among them—would be morally worse than her being required to

choose the only morally acceptable option. Instead, we must differentiate between different types of choices, only some of which, if available, would substantially contribute to agents' well-being. As I will argue later, we can identify the alternatives which we indeed ought to offer to agents by recognizing their role in promoting and reflecting one's agency. *All* actions available to an agent are not such actions, and the confinement objection as it stands gives us no way to differentiate between cases in which confinement poses a moral problem and those in which it is merely an unfortunate consequence of being moral.

It is a fact about the world that we live in that many of the actions we might take would entail performing some moral wrong. Shiffrin and Slote seem to think that this is a problem for highly-confining moral systems. However, I think that this is an unfortunate problem for those of us who seek to act morally, and that the agent, not the moral system, must bear the brunt of the burden here. We shouldn't adjust our moral theories to give agents permissions to perform moral wrongs if they are not already entitled to perform such acts (as I will discuss later, there do seem to be many instances where agents have such entitlements). Instead, agents might seek to identify a wider range of morally good acts within an otherwise-acceptable moral system or they must simply accept the fact that, in some situations, being moral will cause them to have very few permissible options. Gaining a wider range of opportunities at the expense of doing what we otherwise find to be morally wrong is putting the cart before the horse. Although I agree with the claim that agency is not properly recognized by extremist moral systems, I will later argue that only certain exercises of agency can ground a defensible limit to excessive moral demands—there may still be, on my account, situations in which we have

no choice but to do the right thing. Our search for a justification for an objection to highly demanding moral theories continues.

In *The Limits of Morality*, Shelly Kagan offers (and ultimately rejects) two quite plausible justifications for the appeal to cost. Both arguments view extremist theories as failing to account for certain considerations that are of moral significance, which leads extremism to make its costly demands. Kagan argues that in order for the moderate (this is his name for the Defender) to justify the appeal to cost, she must make the (quite plausible) argument that morality must “adequately reflect the nature of persons;” in particular, Kagan is referring to persons’ natural tendency to view their own pleasure and suffering as more salient than those of others.¹⁸ The two arguments reflect two attitudes—either a “negative” stance or a “positive” stance—toward this natural tendency, but both proceed by arguing that morality must recognize this tendency, thereby grounding the appeal to cost. Ultimately, Kagan concludes that neither approach succeeds in justifying the appeal to cost, as both fail to point to any consideration not adequately reflected by a fully objective, impersonal consequentialism.

The first approach Kagan considers he labels the “negative argument.” It takes an ultimately negative view of agents’ natural tendency to favor their own interests, viewing it as an “unfortunate hindrance” that must be accommodated by a moral system—and concludes that doing so requires the granting of options to pursue one’s interests. According to the negative argument, there is a motivational condition on morality that is not met by a moral system that fails to include agent-centered options: for agents to be under a moral requirement, they must be capable of being *motivated by* it, and agents are *not* capable of being motivated to comply with extremely costly moral demands. The

¹⁸ Kagan (1989), p.262

negative argument thus concludes that the pro tanto reason must fail to generate a moral requirement when doing so would result in great cost to the agent: they must be granted *options* in such cases. The premise that is most important (and also most controversial) for the negative argument is that agents' natural bias prevents them from being motivated to act in accordance with extremism in cases of great cost to themselves; to accommodate this motivational shortcoming, the argument concludes that agents must be granted options not to pursue the good when doing so is particularly costly—otherwise, we are left with a moral system which agents are motivationally incapable of adhering to.

Kagan, however, argues that extremist moral systems do meet this motivational requirement. He observes that we typically fail to fully comprehend the implications of our beliefs about others' suffering—as he puts it, our beliefs about their suffering are “pale” instead of being “vivid”, as they are in cases of our own suffering. If we can recognize this, Kagan argues, we can also recognize that our motivations would be different were we to have vivid beliefs about the suffering of others, and can therefore recognize that it is only our *cognitive* shortcomings that prevent us from acting morally.¹⁹ Kagan argues that we can easily see that the more vivid our beliefs become, the more closely our actions would tend to conform to those required by a fully objective, impersonal moral system. Even if individuals are not capable of making all their beliefs vivid, they can still recognize how they would act *were* their beliefs vivid—and this is all Kagan needs, he believes, to show that extremism *does* minimally reflect the nature of persons by satisfying the motivational condition: the only reason we are not motivated to act in accordance with the demands of an extremist moral system is that our beliefs are not sufficiently vivid. The very realization of how I would behave differently with vivid

¹⁹ See *The Limits of Morality*, “Vivid Belief”, 293-281.

beliefs about the suffering of malnourished children in Rajasthan province in India would be sufficient to motivate me to act in accordance with a highly demanding moral system. Therefore, Kagan concludes, I *can* be motivated to aid them, even if I am not *in fact* motivated to do so because my beliefs are currently pale and not vivid. Kagan's key claim is that he only need prove that agents *could* be motivated to act morally, not that they are motivated in their current state.²⁰

Some philosophers, most notably Michael Bratman, have questioned whether Kagan does indeed show that the motivational condition is met for all agents, but I will not take up such arguments here, as I find Kagan's second argument more compelling and highly relevant to the arguments I will develop in the following sections.²¹ Without the negative argument, the only way that Kagan sees for the Defender to escape extremism and ground a defense of the appeal to cost is by pointing to some source of positive moral value reflected by the appeal to cost, and not reflected by extremism. The "positive argument", named for its view about the nature of persons as something of positive moral value, seeks to justify the appeal to cost in this way. Of course, Kagan believes that this argument, too, ultimately fails, and I will discuss his response to it here. Specifically, I will question whether Kagan really constructs the positive argument in its strongest possible form, and whether the positive argument might actually point to some deeper source of value that extremism does indeed fail to account for, providing a justification for the appeal to cost. In the section VI, I will consider arguments by Williams that resemble the positive argument construed in its more plausible form.

²⁰ Kagan does offer some other arguments against the negative argument; all, however, are secondary to his "vivid beliefs" response. However, regardless of their ability to defeat the negative argument, Kagan still offers the Defender another way out of extremism—the positive argument—and this, I believe, is all that is needed to justify certain objections to extremist moral systems.

²¹ See Bratman, "Kagan on the Appeal to Cost".

Two claims form the foundation of the positive argument, according to Kagan. The first, and fundamental, claim is that the nature of persons, and specifically the “personal point of view”, is a source of positive moral value. As a consequence, Kagan concludes that a second feature also holds: that the positive argument must require a moral system to *fully* reflect the nature of persons in order to fully reflect the value of agents’ personal point of view. As Kagan sees it, the positive argument so construed insists so strongly upon the value of this personal point of view that “the nature of persons itself largely forms the ends of morality”.²² This second claim is quite strong—far too strong, I believe, to be an accurate portrayal of the Defender’s more moderate view. I will argue that this second feature is not necessary to the positive argument, and that the Defender can insist that the personal point of view is indeed a source of positive moral value without arguing that this point of view must fully—or even largely—form the ends of morality: to do so would be to insist that this is the only source of moral value, or at least the most important, and the Defender is not committed to this view. Furthermore, the Defender need not adopt such a stance in order to gain agent-centered options, as I will demonstrate later. For now, though, I present Kagan’s objection as it stands.

Kagan locates the dispute between the Defender making the positive argument, and the extremist rejecting it, as one about the “moral point of view”, which he defines as the point of view that takes into account all morally relevant reasons. According to the extremist, the moral point of view is identical to the objective point of view, which “takes into account only those reasons generated by something’s objective value—i.e., only the

²² *The Limits of Morality*, 333. This matches my earlier definition of extremism; Kagan, however, takes this to require a strictly consequentialist approach, while I believe that extremism is possible in other forms.

agent-neutral, pro tanto reason to promote the good”.²³ According to the Defender, however, the moral point of view is distinct from the objective point of view: the former includes reasons not recognized by the latter. Specifically, according to Kagan, the Defender requires that the moral point of view recognize *subjective, personal* reasons as well as objective, impersonal ones. Kagan argues that the Defender is unable to advance any effective argument to this end. He claims that the Defender fails to identify any source of positive moral value that is not already recognized by extremism to the greatest possible extent. Kagan argues that the objective standpoint *already* recognizes reasons given by agents’ subjective points of view, and that the recognition given by extremism is adequate, and indeed the only coherent way of accounting for the competing interests of all agents. So, under Kagan’s construal of the positive argument as one that demands that morality *fully* reflect the nature of persons, the positive argument fails: it is unable to point to a source of moral value that is not recognized by the objective point of view to the fullest extent possible.

It seems, however, that Kagan’s defeat of the positive argument is a hollow one. Specifically, I am troubled by Kagan’s construal of the positive argument as claiming that in order to fully reflect the nature of persons, the moral point of view must *fully* adopt the personal point of view. Instead, it seems that all that is necessary is that the positive argument claim that morality must merely *recognize* the nature of persons, and that morality need not be *entirely* shaped by this nature in order to recognize it. The moral point of view can simultaneously recognize the value of personal and impersonal considerations in its judgments: it only need consider the reasons generated by both the personal and impersonal points of view in its weighing of the overall balance of moral

²³ Kagan, *The Limits of Morality*, 351.

reasons. Morality can acknowledge the significance of agents' subjective attachments to various pursuits while also acknowledging that these attachments, if they are to be morally acceptable, must be pursued in a way that also reflects other, impartial considerations. Reflecting the value of these subjective attachments does not preclude the possibility that these attachments are not the *only* source of value, and that, in particular, impartial considerations will limit the extent to which agents can morally pursue these attachments.

So, the positive argument isn't necessarily as extreme as Kagan construes it to be (although it could be), but is strongest as a slightly more ambitious version of the negative argument, in a sense: like the negative argument, it argues that morality must at least partially accommodate the nature of persons, but unlike the negative argument, it argues that the nature of persons is a source of positive moral value in order to reach this conclusion. The Defender need not insist upon the value of the personal point of view as the primary or ultimate source of moral value; it is only *a* source of moral value that can be considered alongside, and even eclipsed by, other values. So, it seems that Kagan's refutation of the appeal to cost is not as conclusive as he takes it to be. Indeed, I believe that Williams and Waldron, among others, have offered arguments which might be taken as a plausible assertion of the positive moral value of the nature of persons, and as evidence that certain sacrifices required by extremism may indeed be morally impermissible. I will discuss this in sections VI-VIII.

Having explored a variety of options about what exactly the appeal to cost might be founded on, we remain exactly where we started: with a plausible hunch about the appeal to cost and the overdemandingness of extremist moral theories, but a lack of

understanding about what the problem with demanding moral theories really is. We have explored possible problems with the moral theories that produce these excessive demands—that they require agents to shoulder more than their fair share, that they neglect an intrinsic moral fact, that they alienate agents, and others—but to no avail. In terms of cost itself, it appears that there really is no objection that can be mounted to even extremely demanding moral theories; indeed, I will argue that to believe that we can object to a moral theory strictly because it demands a great deal of us is to misunderstand the moral project which generates these demands. Given this failure, the Defender must either concede that there is in fact no problem with demandingness that is not already accounted for by extremism, or the Defender must search for justification for the appeal to cost elsewhere—she must, I argue, point to a flaw in moral theories which make these costly demands. This marks a transition from seeing the demandingness itself as an objectionable feature to seeing the demandingness as a mere *symptom* of a deeper flaw with a moral system that is so demanding.

I will argue for this claim that the problem highlighted by the appeal to cost is *not* actually a problem with moral theories in virtue of their demandingness alone. Instead, I suggest that the appeal to cost is not motivated by an external problem, one requiring us to *limit* morality itself, but instead one internal to morality, which points to some *other* source of moral value that only limits our moral obligation to aid others, not morality itself. The appeal to cost, I will argue, can only be successful if it is able to point such a moral consideration that is left out by highly demanding theories.

V. Unlimiting Moral Requirement

Given the failure of a wide variety of approaches to limit the demandingness of extremist moral theories, we are forced to take seriously the possibility that there is nothing inherently problematic about extreme demands to aid others. As I will argue here, there are few or no limits to what we can be morally called to do, at least in the sense in which we typically speak of moral theories as being “too demanding.” Instead, moral demands can only be excessive if they demand something that is of greater moral value than that which would be achieved by a given action. Because demandingness itself is not an independent problem, we can only look to problems with the demands themselves—do they take all relevant moral reasons into consideration, or do they neglect certain values and therefore make incorrect—and perhaps extremely burdensome—demands? If our moral system itself is acceptable, and if this system has been properly applied to the action under consideration, then we must accept its demands—if we take proper account of the necessary moral considerations, there is nothing *else* for us to consider when determining moral demands, no independent, external limit on moral requirement. When excessive demands are relevant, it will be because their costs can be accounted for in terms of some other moral consideration—and in the following sections, I will turn to an approach that is able to indicate such values, and thus point to considerations which will, in some cases, limit our moral duty *to aid others*, but does not, strictly speaking, limit morality itself.

I am opposing the claim that there is any general limit, relative or absolute, on *how great* the costs imposed by moral demands can become before they are simply unacceptable. This is, at least in part, what motivates Murphy to *dissolve* the problem of

demandingness: it doesn't make sense, given the function of morality (which is, I will argue, to make demands) to say that there exists some point at which morality just demands "too much;" the acceptable level of demandingness depends on what is, morally speaking, at stake. This is one important lesson we learn from the failure of the straightforward justification, which failed to point to any *moral* value to oppose excessive moral demands. If the balance of moral reasons does in fact weigh in favor of an action, then the mere fact that it requires us to make large sacrifices does not generate any *further* reason against it, unless the things being sacrificed are somehow shown to be of moral relevance. In cases of extreme demandingness, the Defender cannot point to the mere quantity of her own resources being sacrificed; the extremist would simply point to the stronger reasons in favor of performing the action in question—this is why the act was morally required, after all.

If the acceptable level of demandingness always depends on what else is at stake, then Judith Lichtenberg is right that there can be "no *general* answer" to the question about where to set the bar for the demandingness of a moral theory: in general, morality's demands have no limit, and we must determine which demands are acceptable as individual cases arise, through the normal means of moral judgment.²⁴ If, despite great costs to an agent, she is morally required to perform some act, then it is certainly unfortunate for her, but the mere fact that she is required to give up a great deal does not generate any *further* reason against her so acting.²⁵ If extremely costly demands

²⁴ "Negative Duties, Positive Duties, and the 'New Harms'" 29.

²⁵ In her discussions of Kantian moral judgment (see "The Practice of Moral Judgment", "Integrity and Impartiality"), Barbara Herman considers such an objection based on demandingness, arguing that agents who face great demands are truly unfortunate, but do not have any legitimate objection to an otherwise-acceptable moral theory. She similarly rejects the claim that those unfortunate agents who have great demands imposed upon them can raise any legitimate objection to morality as a mere consequence of the intrusive, but legitimate, demands morality makes. Scheffler (1986) also concludes that objections to

generated a special reason against requiring agents to perform them, we might also be inclined to think that in these cases, where some extreme moral value would also be brought about (this would have to be so, in order to generate reasons strong enough to outweigh those generated by the agents' own interests), a special reason might be generated in favor of requiring the action. So unless the appeal to cost is dissolved into the identification of a specific moral failure, there is no reason why objections based on extreme moral costs wouldn't be met with arguments based on the analogous extreme moral benefits in such cases.

The conclusion that there is no limit to moral requirement does not mean, however, that our requirements *for beneficence* are unlimited. Instead, it means that if there are considerations that limit our duties to aid others, they must be other moral considerations. I will argue in the following sections that there are some important ways in which extremism's costly demands for beneficence are limited by other sources of moral value—in particular, by ground projects and rights which are central to our agency. The limits on our requirements for beneficence that I will discuss are limits that are *internal* to morality, generated by other sources of moral value, and are not limits on the authority of morality in requiring us to aid others.

The considerations introduced in later sections—based on ground projects and rights—are such moral considerations, values which ought to be, but are not, recognized by impartial extremist systems. Although these agency-based considerations have the effect of reducing our duty to aid others in some cases, their basis is in concepts that form the very foundation of our commitment to morality. The failure of extremist moral

highly-demanding moral theories are best understood as objections to other features of the theory, such as its stringency or overridingness, and not as objections to the demandingness itself.

theories to fully recognize these values is *not* due to an intrinsic problem with excessive demandingness; instead, the failure is one internal to morality, and one of which excessive demands are merely a symptom. The true problem with the demands of extremism, I will argue, is its failure to recognize the moral significance of agency, which requires the addition of a constraint—something the extremist approach refuses to incorporate.

Robert Goodin has offered some particularly illuminating thoughts on the notion of moral demandingness, arguing that demandingness itself presents no problem for a moral system.²⁶ He compares morality to a very demanding person, “a real shrew,” but one who only demands what he is due. What, Goodin asks, can we really complain about in such a case? Goodin makes the fair claim that “a morality demanding only what is morally due can hardly be castigated for that. If anything, we ought be blaming people for not doing the right thing without moral nagging”.²⁷ The problem of demandingness, I argue, is not a problem of silencing the shrew or getting him to nag us less, but one of making sure what the shrew says is correct—that the shrew’s demands rest on an understanding of all relevant considerations.

Goodin also recognizes a rarely mentioned underside to the problem of demandingness: we would see a problem with a moral system that didn’t demand *enough* of agents. A moral system that simply issued recommendations or suggestions, but never actually instructed us to behave in accordance with these recommendations, would seem equally problematic as one that issued extremely costly demands. This suggestion is related to the Defender’s commitment to the existence of the pro tanto reason: the fact

²⁶ Goodin, “Demandingness as a Virtue”.

²⁷ Goodin, “Demandingness as a Virtue”, 2.

that a given action promotes the overall good does indeed seem to be a reason in favor of performing that action, and a moral system which issued no requirements would be one which failed to recognize the *pro tanto* reason—something even the ordinary moralist would object to. Considering such cases in which morality is extremely *undemanding* removes, I believe, some of the intuitive opposition to moral demands that require sacrifices: these cases reveal that we do *want* morality to make demands. We not only believe that a moral system can demand too *much* of us, but that it can also require too *little*—the challenge is finding a moral system that is free of both flaws. We *want* the shrew to speak—we just want it to make the right demands, and a shrew which consistently demands great sacrifices just seems to have gotten something wrong.

The discussion in section IV, as well as Goodin's analysis, force us to confront the possibility that there is no legitimate objection to demanding theories in terms of the demandingness itself, without reference to sacrifices of other sources of moral value. In other words, the right moral system cannot be wrongfully demanding—if there are problems with the demands, they must be due to a flaw in the moral system that makes them, not to some flaw with the demands themselves. The attempts discussed in section IV each failed either because they were unable to point to any moral reason for us to limit the sacrifices imposed by extremism's demands for beneficence (as we saw with the straightforward account) or because the moral values pointed to were not adequate for resisting moral requirements to aid others (the compliance condition, the alienation objection, and the confinement objection). Morality can fail on its own terms—that is, we might mistakenly judge that something is our moral duty when in fact we have failed to

take all necessary considerations into account—but morality cannot make demands which are faulty simply in terms of the sacrifices they require.

Despite my arguments, some intuitive opposition to the idea that morality's demands are potentially unlimited may still persist. I hope that this idea will gain an even stronger appeal later, when I introduce and argue for the moral relevance of some considerations that have been neglected by most contemporary extremist moral systems, but without which the moral project itself simply doesn't make sense. The idea that morality *can* demand a great deal of agents does not mean that it always should or will make such demands—and these other considerations will, in many (but certainly not all) cases, prevent morality from demanding extreme sacrifices. I believe that it is the neglect of these considerations that prompts much of the opposition to the theories that do so, and when they are recognized by our moral judgments, hopefully the notion of unlimited moral requirement will seem more acceptable.

By concluding that moral requirement is in a certain sense unlimited (i.e., that there is no external or general limit to demandingness and that only other sources of *moral* value can limit what morality may demand), I do not mean to argue that we are always morally required to do anything that promotes the good, regardless of the cost. Instead, I am arguing that as long as the balance of morally relevant reasons weigh in favor of an action, then one indeed ought to do it—and that the appeal to cost should be understood as making some claim about how highly demanding theories have failed to properly weigh the moral reasons in a given situation. The pro tanto reason to promote the good is not, as I will argue, the only reason we need to consider in our moral

judgments, and so the act which best promotes the impartial good will not always be morally required.

In light of the fact that morality may ask us to do quite a lot, and that performing these actions may require a great deal of costly sacrifice, we would do well to act to minimize the potential costs of compliance with our duties. Indeed, this is not only a matter of convenience or economics, but might also be construed itself as a moral duty—in order to promote the good as much as is possible with our limited resources, we ought to make promoting it as easy and efficient as possible. Reforming social and political institutions to encourage and normalize (or even require) forms of behavior which promote the good and forbid clear instances of undermining it, developing a moral culture, and minimizing the effort necessary to make an informed contribution to organizations which do effective work to promote the good are all steps to which we should give greater priority. As the true scope of our moral duty becomes clear, I have little doubt that steps will be taken to minimize these costs, but doing so seems especially important as this would encourage moral action by those for whom, without such social and institutional structures in place, moral action seems too daunting and who, as a result, do little or nothing.²⁸

²⁸ Lichtenberg mentions the importance of such efforts to minimize costs: "Changing background conditions against which people act—through law, public policy, and the changing behavior of others—is an essential ingredient to lowering the costs for individuals to comply with norms. Altering what is available—both in the material sense and in the psychological sense of being salient to consciousness—changes both the social infrastructure and the psychological landscape" (31-32). Railton (1984) also suggests such an approach.

VI. Ground Projects and Moral Demands

In the preceding sections I discussed, and ultimately rejected, the possibility that we can identify any problems with the demands made by impartial moral systems in terms of their demandingness itself. This does not settle, however, whether there is some *other* problem with highly demanding moral theories—indeed, there seems to be *something* amiss with them. I suggested that if there is a problem, it must be a problem with the theories that generate these demands. In this section, I will turn to Bernard Williams’ arguments that the impartial point of view adopted by extremist moral systems leaves out an important source of moral value expressed by agents’ “ground projects.” I will argue that Williams’ claims are a form of the positive argument, as they point to a group of important moral considerations which cannot be given their proper recognition from the impartial point of view. Extremist moral theories, then, generate excessive demands because they fail to consider the value of these ground projects—a value which, I will later argue, is part of a larger class of values generated by our agency.

It has become quite common for opponents of extremist moral theories, and of consequentialism in particular, to criticize their failure to account for projects and values that give meaning and shape to an individual’s life, that make it “worth living” in an important sense, beyond their mere status as pleasure-producers or generators of similar positive outcomes. Williams has forcefully critiqued the inability of consequentialism to account for fundamental elements of persons’ character, including the “ground projects” which “are closely related to his existence and which to a significant degree give meaning to his life....providing the motive force which propels him into the future, and gives him

a reason for living”.²⁹ Williams holds that the inability to recognize the moral significance of such projects constitutes a failure to account for the full spectrum of moral reasons, and that this failure actually undermines agents’ very commitment to their own existence—and thus also to their acting morally. In a similar vein, Samuel Scheffler argues that impartial moral systems like consequentialism are, motivationally and otherwise, not “fit for humans” and must be modified to reflect and balance strictly impartial considerations with agents’ own assessments of value, which give their own projects and concerns much greater weight.³⁰ Related arguments have been made by other philosophers, who have highlighted the need for extremist moral systems to take account of “personal” considerations.³¹

Williams’ objections in particular can be understood as a version of the positive argument discussed above. Williams holds that there is an important source of moral value—agents’ *character*, which is reflected in their ground projects—that cannot be recognized from the impartial point of view adopted by extremist theories like Kagan’s. These impartial theories require that we abstract away from all features of a situation except insofar as they can be treated as features of any morally similar situation. In this abstraction, Williams argues, these theories ignore elements of our character which are of moral relevance—they fail to recognize that my relation to my own projects is *not* the same as my relation to others’ projects. This is problematic for two reasons: first, my ability to pursue these projects expresses the intrinsic value of my agency. Second,

²⁹ Williams, “Persons, character and morality,” 12-13. See also “Integrity” in Smart and Williams, *Utilitarianism: For and Against*. Williams also criticizes the inability of impartial moral systems, utilitarianism in particular, to account for the separateness of persons—the importance of *who* performs an action—but I do not take on this objection here, as I find his discussion of ground projects more relevant and more illuminating for discussion of the limitations of moral requirement.

³⁰ See Scheffler, *Human Morality*.

³¹ See Hurley (1992), and Portmore (2000).

because my agency has instrumental value in allowing me to act morally, this abstraction risks detaching me from my own life and thus from morality itself: my own projects, he says, “are the condition of my existence, in the sense that unless I am propelled forward by the conatus of desire, project and interest, it is unclear why I should go on at all”.³² My projects are what give my life meaning; they are the reason I never need to ask myself why I go on at all. If I am forced to sacrifice these projects, my reason for living may be taken away—and with it, my reason for being moral. This relation between my projects and the meaning of my life is unique to my own projects and, if violated, undermines my ability to promote the projects of others at all. Williams thus sees extremist moral theories as not only making a mistake in judgment by failing to recognize agents’ important relation to their own projects, but as making a mistake which ultimately undermines’ agents’ connection to their own life and to morality itself.

Williams argues that in order to fix this blind spot in impartial theories, agents’ ground projects ought to be given *some* consideration when determining what demands can be made of them—both because of the intrinsic moral value expressed by my projects, and because failing to consider this special relation between an agent and her projects is self-undermining: moral requirements which erode my commitment to living my life also, therefore, erode my commitment to morality. He does not, however, insist that these personal considerations be allowed to independently constitute the moral point of view—which was Kagan’s primary concern with the positive argument. Williams merely insists that morality acknowledge and be compatible with the value that agents’ ground projects express. His argument thus avoids Kagan’s criticism that the positive argument must allow the personal point of view (the point of view from which a given agent reasons) to

³² “Persons, character and morality,” 12.

fully constitute the moral point of view (the point of view adopted by a moral judge).

Williams simply holds that in order to fully accommodate the moral value of agents' ground projects, the moral point of view must at least *recognize* such considerations by including them in the balance of other reasons that determine our obligations. Failing to do so would be to operate on "a misrepresentation, since it leaves out what both limits and helps to define that aspect of thought".³³

Williams also believes that the impartial point of view is fundamentally incapable of recognizing such considerations. He rightly points out that the process of abstraction central to strictly impartial moral systems inevitably produces an "impoverished and abstract" view of persons, "omitting their character" by neglecting the ways in which agents are especially attached to their own projects. However, he goes further:

[T]he point is that somewhere...one reaches the necessity that such things as deep attachments to other persons will express themselves in the world in ways which cannot at the same time embody the impartial view, and that they also run the risk of offending against it.

They run that risk if they exist at all; yet unless such things exist, there will not be enough substance or conviction in a man's life to compel his allegiance to life itself. Life has to have substance if it is to have sense, including adherence to the impartial moral system; but if it has substance, then it cannot grant supreme importance to the impartial system, and that system's hold on it will be, at the limit, insecure.³⁴

Strictly impartial moral systems and the agents' own point of view are, Williams believes, incompatible: they will come into conflict whenever morality and our own interests diverge. When this happens, impartial moral systems *must* always prevail—and in doing so, will demand that agents sacrifice the "substance" of their own lives—their character, as reflected in their own ground projects. Williams sees systems like Kagan's, which insist upon the "supreme importance" of an impartial system, as forced to promote

³³ "Persons, character and morality," 19.

³⁴ "Persons, character and morality" 18.

the impartial good at the expense of recognizing the moral value of agents' identity, as expressed by their ground projects. Williams, then, is making an argument for the appeal to cost which holds that excessive demands to aid others are indeed too costly because they fail to recognize the moral significance of agents' connection to themselves—and this neglect of an important aspect of personal identity is, he believes, a moral failure which is inherent to impartial moral systems like extremism, which are fundamentally incapable of recognizing agents' categorical commitment to their ground projects. Williams is thus addressing the concerns raised by the alienation objection about agents' connection to themselves and their own projects being unrecognized by morality. Williams, however, goes further than the alienation objection in giving an account of why this alienation is morally relevant—because it is connected to our identity and to our ability to be moral at all—and in allowing us to distinguish between cases where this alienation is morally problematic (when it breaks these ties) and when it is morally necessary (when it is needed to maintain these ties).

I agree with Williams that agents' ground projects are a highly significant source of moral value, and that the failure to acknowledge this value is a cause of unreasonably costly moral demands. I hope to show, however, that such values are indeed central to the project of even impartial moral theories, and that such theories ought—and *can*—recognize these values. While this flaw is indeed inherent to extremism, which *only* recognizes the value of promoting the overall good, other impartial moral systems can recognize the significance of the “substance” of life, because there is a way of recognizing an impartial feature beneath such values: agency, which I will discuss in detail in the final section.

Before I discuss how the significance of agents' ground projects can be incorporated into even impartial moral systems, I want to point out a tension that is left unresolved by Williams' discussion. He devotes little attention to the delicate balance that morality must achieve between recognizing our own projects and those of others. Our ground projects will conflict with others', and an appropriately demanding morality will be able to acknowledge the value of both sets of ground projects while acknowledging that not all projects can always be permissibly pursued. Even when agents' ground projects are incorporated into the balance of moral reasons, they will be subject to weighing just as other considerations are: this is a consequence of acknowledging that the personal point of view does not alone *constitute* the moral point of view. While ground projects are certainly an important ideal that we should strive to recognize insofar as possible, we live in a world in which our ideals often come into conflict.³⁵ I agree that, absent other relevant moral constraints, all agents should be able to pursue their ground projects and allow those projects to shape their lives as they see fit, free of intrusion. However, it is an unfortunate fact that we usually cannot pursue our own ground projects without committing some moral wrong, or disregarding the projects of another. Indeed, it will often be the case that persons will have mutually exclusive ground projects, and morality will be forced to sacrifice one set. We are then forced to evaluate the relative strengths of the two conflicting considerations—the value of my own projects (including their special value *for me*), and the value of yours. While I agree with Williams that ground projects should be given consideration when determining one's moral obligation, it is unlikely that such evaluation will (or should) always settle in favor

³⁵ Of course, the mere fact that our ideals cannot always become reality is not itself a reason to abandon them. In *Moral Clarity* Susan Neiman emphasizes the importance of striving towards ideals and using them for orientation, as a guide to our actions, even though these ideals will rarely, if ever, be achieved.

of an agent's own ground projects—these are not the only source of value at stake. The problem with impartial moral theories is not that they require an agent to temporarily curtail or abandon her ground projects in order to promote those of others—this would happen even under a moral theory which meets Williams' criticisms, when agents' projects come into conflict. The problem is that impartial moral theories require agents to curtail or abandon their pursuit of their ground projects *because* such theories fail to even recognize their moral significance. Recognizing the moral value of agents' ground projects means recognizing them for *all* agents—and sometimes, as a result, we may be required to sacrifice even our own projects. This must happen if morality is to adequately balance the value of our own projects and the value of others'. I believe that we can have a balanced account of the subjective significance of ground projects as Williams insists, but without necessarily abandoning impartial moral systems altogether—we can employ an impartial moral theory which contains an agency-based constraint. This, I think, is a solution that would satisfy both Williams and Kagan (but still requires abandoning Kagan's extremism).

In refuting the positive argument, Kagan argues that the Defender must point to a source of value which we are unwilling to deny, but which can only be given its full or necessary recognition from the subjective standpoint. While Williams says little about the problems encountered when ground projects come into conflict, his theory must be able to accommodate such situations—they occur with great frequency and are of serious concern, and Kagan argues that the impossibility of recognizing the subjective significance of all agents' conflicting ground projects is exactly why his impartial approach is the only feasible one.

Kagan addresses concerns about extremism's willingness to sacrifice anything under certain circumstances—even ground projects—by pointing out that such sacrifices will always be directed at producing even greater moral benefits, which is the only morally acceptable way of settling these conflicts:

Even if promoting the good should remove various elements of color, leaving the world grey, it is crucially important to recall that the reduction in overall good resulting from adopting a subjective perspective reduces the overall color of the world even more: a small splash of color is only purchased at the cost of plunging far greater areas into dark black.³⁶

Kagan sees the extremist approach as the only way to resolve such problems. The only reason someone is morally required to curtail her pursuit of her own ground projects in order to aid others is that these persons are, quite frequently, in danger of losing the most basic necessities required to form such projects themselves—they are face-to-face with pitch black. Permitting agents to pursue their own ground projects while ignoring the stronger claims generated by others' ground projects would, Kagan holds, be unacceptable—unless the Defender can give some reason why demanding that an agent sacrifice her own ground projects in this way is not actually the morally-best option. However, as I have argued, Williams offers at least two reasons: First, the painter's free choice of colors represents its own source of value, one that cannot be reflected in a painting that is made according to the demands of Kagan's morality, which would resemble a paint-by-numbers system. Furthermore, if we do not allow the painter to color, at least within certain limits, as she sees fit, then she may cease to have any reason to paint at all.

While both Kagan and Williams and Williams agree that impartial moral systems are incapable of giving any recognition to ground projects beyond their purely impartial

³⁶ Kagan, *The Limits of Morality*, 359.

significance, Kagan thinks this is an unfortunate fact for agents' ground projects; Williams thinks this is an unfortunate fact for impartial moral systems. I, however, believe that the failure to recognize ground projects is not inherent to impartial moral systems—but that it is indeed a failure of Kagan's. While many philosophers advocating such impartial conceptions have portrayed them as leaving little space for the moral importance of ground projects, I believe that such systems can indeed accommodate ground projects—and can do so for reasons Williams would accept. The "impoverished and abstract" view of the character of persons that is the target of Williams' criticism is due to a narrow understanding of morality's function and the type of agents it governs, and I believe that even an impartial moral system can recognize ground projects through a constraint.³⁷ I will discuss this constraint in detail later; in what follows I am only interested in showing that impartial moral systems are indeed capable of recognizing the moral relevance of ground projects.

As Williams notes, the moral point of view adopted by impartial moral systems abstracts from particular features of situations "except in so far as these can be treated as universal features of any morally similar situation".³⁸ I see no reason, however, why ground projects could not qualify as sufficiently universal features so as to be recognized and granted weight by an impartial moral system. While the specific content of persons' projects varies, plans which extend over time, shape agents' deliberations and actions, and give meaning to their lives seem to be present in the lives of most, if not all, persons under this general description. These projects seem, then, to be a sufficiently generalizable feature to be recognized from an impartial perspective. If so, then these

³⁷ "Persons, character and morality," 4

³⁸ *Ibid.*, 4.

projects could enter into the balance of reasons considered by such a moral system.³⁹ Indeed, given the force and plausibility of Williams' arguments and the clear significance and generalizability of ground projects, it appears that the ability to account for such commitments might be at least a very large part of the systemic flaw underlying the excessive and objectionable demands of extremist moral systems. As I mention later, even impartial and extremist moral systems are committed to recognizing the moral significance of agency. As Williams makes clear, these ground projects, and their continuity over time, do represent an important aspect of persons' character—and as such, of their agency. So any moral system will, as Williams argues, need to recognize these projects if it is to avoid making demands which inhibit and damage our agency. I simply disagree with Williams that any radical departure from Kantian or consequentialist ethics is necessary to do so; an ethical system which borrows concepts from both might prove entirely sufficient to account for the moral significance of these projects.

Williams' critique does raise some concerns that I will not address here, but that have important implications for questions about moral demands. In particular, Williams raises important questions about the authority of morality—about where our reasons to be moral come from, and whether morality must, ultimately, trump other sources of value. While I will not attempt to answer these questions, I do think that such questions might be just another way of voicing concerns about the limits of moral requirement discussed here. Questions about whether moral obligations are limited by other types of value are linked to questions about whether certain types of moral considerations can limit our

³⁹ My claim that impartial moral systems can account for the reasons generated by ground projects is independent of the question of whether “personal” reasons ought also to be included in moral judgment. If personal reasons ought also to be included, then the moral significance of ground projects seems to follow uncontroversially. The project of arguing for the moral relevance of such reasons is the explicit goal of Hurley (1992).

requirements to aid others.⁴⁰ I approach the question by identifying *moral* considerations which sometimes limit our obligations to aid others, and it is beyond the scope of this essay for me to say much about how questions about the authority of morality might limit moral requirement. Williams' approach also raises concerns about situations where ground projects might come into direct conflict with morality, questioning whether morality really ought to have the last word in such situations. I think this is a very important and interesting problem, but I will not attempt to solve it here. I am only seeking to address questions about the content of morality, not its authority—although I recognize that these are intimately connected—and, in particular, I am interested in demonstrating that the content of morality's demands is compatible with living a full, satisfying human life, which appears to be the primary concern motivating both approaches.

VII. Rights Left Out

Some rights theorists have expressed concerns similar to Williams' which help make clear the significance of agency in limiting our moral obligations to aid others. In particular, a rights-based approach reveals that we do not necessarily have to care deeply about certain projects order for them be morally relevant—even if they fail to have any deep significance *for us*, certain activities are so significant that we have rights to perform them which can limit our moral obligations to aid others. As a result of our

⁴⁰ Susan Wolf, for example, sees the problem as one external to morality—as one which, ultimately, limits the influence that morality should have on our lives (See “Moral Saints”). Under the tripartite conception of “The Art of Life” that Mill adopts, the question is also framed as one of the relation of the moral sphere to other spheres (for a discussion of this, see Jacobson, “Utilitarianism Without Consequentialism: The Case of John Stuart Mill”). However, I see little difference between such an approach and one that instead tries to incorporate concerns about our requirements preventing us from realizing certain sorts of value *into* morality.

possession of certain rights, these philosophers argue, we are permitted to perform acts which would otherwise be morally wrong—that is, we have “a right to do wrong” in certain instances.⁴¹ If rights—including rights to pursue ground projects—are not given their proper recognition by impartial moral systems, then this would explain why such systems generate demands which appear to neglect important sources of value.

The notion of moral rights being capable of resisting moral demands is not a new one.⁴² Here I am interested in Jeremy Waldron’s approach, however, because it directly addresses the problem of acting against morality’s demands, and offers a plausible explanation for why moral requirements to aid others might be subject to an important constraint generated by our rights. On Waldron’s account, no matter how demanding morality may be, in certain cases we can legitimately exercise our rights in violation of these requirements. Indeed, we can assert our rights while recognizing that acting in this way may be “inconsiderate, destructive, wasteful, deceitful, and just plain wrong”.⁴³ Waldron’s stance, I argue, is in many ways a more general version of Williams’, and one which makes clearer the basis for Williams’ claims: just as morality must recognize the significance of ground projects despite strong moral requirements to aid others, it must also accommodate agents’ freedom to act in other ways in which they have rights to act.

In “A Right to Do Wrong,” Waldron argues that agents’ negative rights grant them permissions to perform acts which are morally wrong. He claims that all rights

⁴¹ This is, of course, the title of an article by Jeremy Waldron on this topic.

⁴² Nozick (1974) argues that individual rights severely constrain the structure of morally acceptable political institutions. Kagan (1989) denies that rights can ground any opposition to requirements to promote the good (216-230). In a more contemporary discussion, in “Rights and the Duty to Bring Aid,” John Arthur considers the moral significance of persons’ property rights to be important in resisting moral requirements to aid those affected by global poverty. Perhaps the foundational work outlining our rights to act as we choose, even in ways which may be nonideal or even morally wrong, is of course Mill’s *On Liberty*.

⁴³ “A Right to Do Wrong,” 37.

entail a corresponding moral injunction against interfering, and that, as a result, preventing someone from exercising their rights would have moral consequences above and beyond the moral ramifications of the act in question. So, I am permitted to exercise my rights even when this involves performing actions that are morally wrong, and it would be morally worse for someone to violate this right by interfering with my acting. Waldron notes that rights preserve individual choices in situations where the decisions “are seen as crucial to personal integrity. To make a decision in these areas is to decide what person one is to be”.⁴⁴ Given that there are protected “clusters” of actions in which agents’ freedom to choose is protected, and that not all available actions within these clusters will be morally required or even morally permissible, Waldron concludes that there exist, at least in some circumstances, rights to do wrong. Waldron’s approach suggests that discussions of the demands of morality obscure a certain important fact: that even when we may have overriding moral reason to do something—even when morality may be fully entitled to demand that we do it—our rights still permit us to choose not to, regardless of this demand.

He does not think these rights make such actions morally *right*, but instead argues that rights protect our ability to act in ways which are morally wrong. So although I may be within my rights to refuse to give a stranger in the street the time of day, “we cannot mean to claim that actions of these sorts are sealed off from moral criticism and evaluation”—the act is still wrong.⁴⁵ I do not fully agree with Waldron’s use of the phrase “right to do wrong;” I think that if we have a right to do something, then it is morally permissible for us to act in this way. This is part of why I believe an agency-based

⁴⁴ “A Right to Do Wrong,” 34.

⁴⁵ Waldron, “A Right to Do Wrong,” 37.

constraint is required: morality needs to recognize the moral significance of the acts central to “personal integrity” that both Williams and Waldron draw our attention to. This disagreement, however, is fairly minor, and I do agree with Waldon’s overall point: that reasons generated by agents’ rights can, in at least some instances, resist their moral obligations to aid others.

The parallels between Waldron’s discussion of rights and Williams’ discussion of ground projects and integrity are quite clear. That Waldron’s emphasis on personal integrity and shaping one’s own life sounds just like Williams is no coincidence. Indeed, I believe that Waldron’s arguments simply refer to a general class of considerations—rights—of which rights to pursue ground projects are a particular instance, and I believe that both are rooted in the same moral values: the capacity for choice, attachment, and rational action that I will later discuss as the concept of agency. Just as the ground projects that Williams discusses are central to individuals’ exercise of their agency, the rights that Waldron discusses clearly refer to other, similar “clusters” of actions which “have a special importance for individual integrity and self-constitution”—these actions are activities constitutive of agency.⁴⁶ The significance of these capacities for integrity and self-constitution is rooted in the significance of agency, which I will argue is a necessary value of the moral project. The difference, however, is that for Waldron, these rights are generated *whether or not* agents care about the activities: we have rights to these clusters of actions simply in virtue of their significance for our individuality and identity—a significance which holds whether or not we recognize or care about it.

⁴⁶ *Ibid.*, 34. Among the clusters of such actions, according to Waldron, are “individuals’ political activities, their intimate relations with others, their public expressions of opinion, their choice of associates, their participation in self-governing groups and organizations, particularly political organizations and trade unions, [and] their choice of an occupation” (34-35).

Waldron's account thus grants agents rights to a larger range of actions than Williams' account does, as these rights are not even linked to agents' own conception of their identity.

Like Williams, Waldron points to the tension between impartial moral systems and agents' own lives having value *for the agents*, arguing that if the latter is to be preserved, then agents must be granted rights. Waldron argues that if we were only allowed to exercise our rights to perform actions that morality approved or was indifferent toward, such as "[t]he decision to begin shaving on chin rather than cheek, the choice between strawberry and banana ice cream, the actions of dressing for dinner and avoiding the cracks on the sidewalk," then our rights would only protect acts that were so trivial as to seem meaningless—and our connection to our own lives would be eroded as morality made the most important decisions for us, leaving only the tiniest choices up to the individual.⁴⁷ Our individuality is often expressed in the actions which have a great enough impact on the world to elicit moral judgments, sometimes of moral wrongness. Allowing morality alone to determine which actions we are permitted to choose, and failing to take into account the importance these actions hold for our identity, would dilute and limit the available actions to such an extent that the range of choices presented to us would "lose the link with the *importance* of certain individual decisions which...is crucial in their defense"—we appeal to rights precisely because the actions at stake are *important*, not because they seem insignificant and unlikely to elicit any moral judgment (or a necessarily positive one).⁴⁸

⁴⁷ *Ibid*, 36.

⁴⁸ *Ibid*, 36.

While I agree with Waldron that moral systems must recognize rights to perform actions that might otherwise be morally impermissible, I believe that we must be clearer about the source of these rights—both so that we might begin to identify which acts they apply to, and so that we can modify our moral theories to accommodate them. At this point it will be clear that I see agency as (at least in large part) the basis of these rights. Allowing other moral considerations to constrain decisions which are truly central to an individual’s agency would be self-undermining in precisely the same way that a moral system which required an agent to sacrifice her ground projects would be.⁴⁹ Waldron and Williams are highlighting an important fact about agency which moral systems must, and indeed can, reflect: *agents have rights*. In strictly requiring agents to promote the impartial good, extremism actually infringes on their agency—a source of the value that morality seeks to promote—by infringing upon rights which protect agents’ ability to express this value.

Waldron also points out why an approach like his own or Williams’ is superior to the confinement objection discussed above, and makes clear how this agency-based approach builds on the confinement objection: while the confinement theorist “indiscriminately” holds that *all* decisions are morally significant exercises of individuality and ought to be protected, Waldron and Williams instead hold that “some of the choices that individuals make are more important than others,” and they offer us a way to differentiate between the important and the unimportant choices.⁵⁰ The confinement objection fails because it holds that *all* decisions an agent makes represent

⁴⁹ Again, however, I remind the reader that I do not discuss cases of inherently immoral ground projects here—I take for granted that most agents’ ground projects are acceptable to, and even formed within, the demands of morality.

⁵⁰ *Ibid*, 34.

morally-relevant exercises of agency which ought to be protected, and thus that all moral demands which restrict these choices are unacceptable. Waldron, however, rightly holds that only *certain* choices are protected by this right—those central to our identity or self-constitution—while others are indeed subject to moral requirement and confinement.

Morality cannot be permitted to limit our options so much that our decisions fail to even be meaningful ones; however, we also must accept that our ability to choose must be in a delicate equilibrium with other moral considerations, like the needs and projects of others.

This approach also seems to capture what was appealing about the alienation objection without having the same failings. While we recognized that alienation does typically seem to be an unfavorable feature of a moral demand, I also mentioned that alienation is sometimes crucial to *autonomy*—we must undergo it in order to remain autonomous. Autonomy is, clearly, a necessary capacity for rational agency—and the agency-based account I am developing here is able to distinguish between cases in which this alienation is permissible and even morally desirable—i.e., when it is necessary to our autonomy and therefore our agency, or when other moral considerations warrant such alienation—and cases in which the alienation is truly problematic. The concerns raised by both the confinement objection and the alienation objection, then, are addressed by an agency-based approach, which is able to discriminate between morally-relevant instances of confinement or alienation and those in which such confinement or alienation is morally necessary.

VIII. A Morality Fit for Agents

My arguments in the previous sections have made clear that, while I believe that there is no basis for an external limitation on morality's demands, I do believe that extremist moral systems fail to recognize important sources of moral value, and, as a result, make highly—and incorrectly—demanding judgments. In this section, I will outline in detail an agency-based constraint which I take to be necessary to repair impartial moral systems like extremism. I will argue that we must recognize that in order to have any commitment to or participate in the moral project, we must be *agents*, and so a morality must preserve each individuals' capacity for agency, never requiring her to forego her own agency in order to act morally. Acknowledging this limitation on our requirement to aid others will help address concerns that acknowledging any duty of beneficence can open up the floodgates for unlimited, intrusive, and seemingly unbearable moral demands which can peck away at our lives.⁵¹

Before I discuss this agency-based constraint in greater detail, I owe the reader an account of what agency is, why it is of such moral significance, and why this moral significance entails the special protection of a constraint. I will offer a general description of agency itself—but by no means an exhaustive one—before offering two arguments for the moral significance of agency. First, I will argue that positing the moral significance of agency is necessary to explain a wide variety of moral judgments we wish to make. Second, I will argue that the moral significance of agency lies at the basis of our moral obligations to aid others, and that it ought to be equally significant in limiting these

⁵¹ Barbara Herman discusses this issue, among others, in “The Scope of Moral Requirement”, and much of my thinking on this issue has been guided by her penetrating insights into the problem. Lichtenberg also discusses the importance of finding an acceptable level of moral demands in “Negative Duties, Positive Duties, and the “New Harms””.

obligations—indeed, this account is attractive precisely because it points to a single source of value as underlying both our moral obligations and their limits. Finally, I will argue that the only way that the moral value of agency can be properly recognized is by an “agent-relative” constraint; one which, necessarily, is inaccessible to extremist moral theories like Kagan’s (which deny that *any* constraints or options are permissible), but which is compatible with an otherwise-impartial moral system.

While giving an exhaustive account of agency is beyond the scope of this paper, I will identify some aspects of agency which are necessary, but perhaps not sufficient, conditions. My discussion of the moral significance of agency only requires such a basic analysis, although a more detailed analysis would be required to give any account of the precise boundaries of moral requirements to aid others (which, as I stated, I will not attempt to give).

The word ‘agent’ is, of course, etymologically and conceptually linked to the (similarly nebulous) concept of action. Putting aside some of the debate about how precisely to define action, we can recognize that being an agent requires certain capacities that are also required in order to *act*, in the sense of intentional action. Among these are the capacities for rational thought, forming and pursuing a conception of the good and other plans and attachments over time, and the ability to reflect on the desirability of one’s own desires, projects, and ends. I simply take for granted here that these are features of agency, but I believe that there are good reasons for accepting them—one need only reflect on what one’s life would be like without any one of them, and whether one could truly “act” without them. An agent, then, possesses these capacities—and, insofar as

we wish to preserve our agency, we must preserve these capacities in those who have them, and cultivate them in those who do not.

My first argument that agency is a concept of *moral* relevance is an explanatory one: agency must be a moral value, and an important one, if a wide variety of our moral judgments are to be explained. Agency, for example, figures deeply in plausible objections like the alienation objection discussed above—when someone is alienated from her own plans and attachments, her agency is at stake (as I argued, however, sometimes such alienation is actually *necessary* in order to be an autonomous agent—this is why the objection failed). The confinement objection, which also holds some intuitive plausibility, also appeals to agency: If I am not offered a range of morally-acceptable options to choose from, and instead have my actions dictated to me by morality, then I am unable to exercise my agency in ways that are morally permissible (again, however, this objection failed). The notion that agents’ ground projects are morally significant also relies on the assumption that agency is morally significant—the mere fact that these agents choose to pursue these projects, that they deem them good or desirable (these are all exercises of agency) is what gives the projects moral worth on this account. Our beliefs that every person is entitled to education or that even criminals ought not suffer unnecessary or cruel punishment reflect a belief in a sort of moral value that, without the concept of agency (and the privileges that come with it, like rights), seems extremely difficult to justify or explain.

The second argument for the moral significance of agency seeks to understand morality “from the ground up”—that is, by reflecting on the basis of our moral obligations as a whole. Whatever other factors form the basis of our moral obligations, our agency

clearly must be one of them. While their specific understandings of agency and the content of morality vary, both Kantians and Millians broadly agree on this point: our ability to form attachments, reflect on them, and express them through our actions is essential to the value of human life. It is, at least in part, this value of human life—of our *agency*—which underlies our obligations to others: when people are at risk of losing their right to freely express themselves, to have a say in how they are governed, or, to take the extreme case, when they are at risk of losing their life, each of these constitutes a loss of *agency*, regardless of the specific circumstances. The *limit* on our obligation to aid others is also rooted in agency—our own. The fact that an appeal to agency lies beneath both the basis of our moral obligation to others (we take their agency to be something of intrinsic moral worth) as well as the limits of this obligation (we also take it that we cannot be required to sacrifice our own agency in order to aid others) is part of what makes an agency-based account so appealing. Barbara Herman, following Kant, sees the two obligatory ends—our own perfection and the happiness of others, which generate and limit our obligations—as “the complete material specification of rational nature as an end in itself for human rational agents”. These two ends, which she argues are the *only* two obligatory moral ends, are both rooted in the same value—rational agency.⁵² This is what allows the agency-based limitation on moral demands to aid others avoid the ad hoc character of many other attempts to limit such demands—it uses the source of our moral duty to aid others to also explain the limits of this duty.

Given this moral significance of agency, the extremist might wonder how the moral value of agency could prevent moral demands in cases where, in giving up certain projects central to her agency, someone might promote the agency of multiple other

⁵² Herman, “The Scope of Moral Requirement,” 240.

persons. The answer is related to Williams' observation that when an agent's own projects are undermined, her very connection with her own life—and with it her interest in being around in the world and living morally—is removed. The respecting of her own agency, then, is prior to her obligation to aid others, in the sense that if she sacrifices the conditions of her agency in the process, her ability to act morally in the future is at risk of disappearing along with her ground projects—without her ground projects, she will lack one of the necessary features of agency. This undermines her ability to act morally at all: the capacities which constitute agency—including rational deliberation and having attachments to others (which enable one to understand others' attachments)—are necessary to do so. I noted that agency has both intrinsic *and* instrumental value, and a constraint against action which violates this agency is necessary to express both. The intrinsic value of agency fails to be expressed unless agents' deep attachments to their projects are recognized as more than mere positive outcomes (however these are conceived under various forms of extremist consequentialism), as this fails to allow agents to develop these attachments in a way that demonstrates the projects' significance for them. The instrumental value of agency also requires a constraint to protect it, because this agency is necessary for us to act morally at all. If morality were to demand that we sacrifice our agency—which is both the source of our obligations *and* necessary to carry out those obligations—then we would cease to have reason (or the means) to act morally. Extremism's demands for any agent to sacrifice her own agency would, thus, be self-undermining, diminishing something of intrinsic moral value and detaching the person sacrificed from the moral project altogether.

Nagel frames this discussion in terms of “agent-neutral” and “agent-relative” principles.⁵³ Agent-neutral principles are principles without reference to particular agents, and which refer only to features of a situation that are morally relevant and sufficiently generalizable to all morally similar situations. These are the principles which impartial moral systems recognize. Agent-relative principles, however, apply to specific agents and are not generalizable: instead of assigning value to particular states of affairs, which is what agent-neutral principles do, agent-relative principles require specific agents not to act in a certain way, without stating anything further about what she must do to prevent others from acting in this way. Agent-relative constraints hold even when, by acting against them, the agent might be able to prevent multiple identical violations of the constraint. The constraint which must be adopted to recognize agency, I am arguing, is an agent-relative one: it holds that *each* agent is required to respect the moral significance of agency *herself*, and that, in particular, she cannot be required to violate her own agency or to sacrifice the conditions necessary to sustain it in others. This is the only way to make sense of the moral value of agency; it is a value which, if agents are required to violate it, undermines their own integrity, their connection to their own lives and to morality itself.

This distinction, then, makes clear how this version of the positive argument would answer Kagan’s replies: the very moral value which, at least in part, underlies our moral obligation to aid others—agency—is the same moral value which can only be recognized via an agent-relative constraint, one which holds that it is always impermissible for an agent to be morally required to violate the conditions of her own agency. If Kagan were to insist that morality can demand that agents sacrifice their own

⁵³ See Nagel, “Personal Rights and Public Space,” especially p. 34-40.

agency to promote that of others, he would be failing to recognize the intrinsic value of agency which lies at the foundation of morality's demands as well as undermining the instrumentally-valuable capacities which allow agents to be moral. Because agency itself is a fundamental moral value, extremism's reallocating of resources from one agent to even several others would have moral ramifications beyond the mere reallocating of resources—it would be to violate that individual's agency and to prevent her from acting morally in the future—an unacceptable outcome even on the extremist's own terms, for the extremist must, I have argued, appeal to the moral significance of agency if her judgments are to be explained.

Kagan's extremism will undoubtedly make demands which undermine agents' commitments to their own lives, eroding their connection to morality and their will to live at all. Indeed, Kagan himself admits this: "Requiring that an agent promote the [impartial] good effectively cuts off the possibility of his molding his own life. It demands that he be ready to sacrifice all that to which (and all those to whom) he has devoted himself".⁵⁴ Kagan, however, thinks that there is no morally acceptable way to avoid this unfortunate feature of impartial morality. We now have an account of exactly why such demands are problematic for extremism, and how this unacceptable feature can be avoided. The only way that extremism can save itself from this pitfall is to recognize the value, for agents, of the ground projects to which they are "categorically" committed. The only way that these projects can be fully protected and be recognized as valuable *for the reasons that agents value them* is by incorporating an agent-relative constraint, which categorically protects agency, into our moral theories.

⁵⁴ *The Limits of Morality*, 233.

As I mentioned above, my agreement with Williams ends at his claim that all impartial moral systems are *incapable* of accounting for the moral significance of the core aspects of agency. I see “moral philosophy’s habit...of treating persons in abstraction from character” as nothing more than that—a habit, and one that an agency-based constraint can rectify. As I argued, impartial theories are capable of acknowledging the significance of agency through a constraint—but the adoption of a constraint means that we must reject extremism. Barbara Herman notes that Williams’ concerns about the inability of impartial moral systems to account for the value *for agents* of their ground projects applies only to pure consequentialist theories like extremism, which, even when they allow agents to pursue their ground projects, only do so on the condition that they are the utility-maximizing course of action in the situation.⁵⁵ However, a morality which emphasizes the moral significance of agency is able to “respect the agent’s attachments to his projects in a way that permits his actions to be the expression of those attachments”—indeed, it is the *fact that* these actions are expressions of *his* agency which permits them, on my account.⁵⁶ Even an impartial consequentialist system, supplemented by this agency-based constraint could, it seems, recognize the value of agents’ projects *for them* while also otherwise maintaining the impartiality that is essential to moral judgment. On this account, our moral obligations to aid others are limited in that these obligations cannot interfere with agency or the projects central to it, and they cannot do so *because* of the value of these projects for the agent. By contrast, a classical utilitarian or an extremist like Kagan would only be able to see this constraint in terms of maximizing moral value; doing so, however, obscures *why* these projects are important, and the fact that they can

⁵⁵ Herman, “Integrity and Impartiality,” in *The Practice of Moral Judgment*. See especially Section III.

⁵⁶ Herman, “Integrity and Impartiality,” 39.

be more important for an agent herself than they might be on a strictly impartial assessment.

The important thing that even a moral system with an agency-based constraint must recognize, though, is that pursuing some ground projects will be morally impermissible, at least in some circumstances—despite the agent’s connection to them. This need not, however, be seen as a failure of such moral systems to account for an agent’s integrity:

[A]t the limit, where conflict with morality is serious and unavoidable, morality must win. The “victory” of morality does not diminish the value of the projects (unless it is the project itself that is judged immoral), and so does not constitute an attack on the agent’s integrity. It is rather an indication that in this world, or in these circumstances, there is no permissible way to continue this project.⁵⁷

While morality must be shaped by agents’ projects, their pursuit of these projects must also be shaped by morality; the two must be in equilibrium. Agents who are committed to acting morally will have to recognize that sometimes, even when their integrity is taken into account, that morality cannot endorse these projects—in such cases, there is no way to resolve the conflict of ideals. This does not mean, however, that their agency will be thus eliminated—their commitment to morality itself is a ground project, and one which, in such cases, will precede their commitment to other projects. Thus, while a moral system which includes requirements for beneficence is indeed compatible with agents living a “full” life of attachments, commitments and projects, such a system will still, at times, require sacrifices—but these sacrifices will be in order to further a project to which they are committed, namely, morality. The demands made by a moral theory with an agency-based constraint might still be quite significant in certain situations—and when the theory

⁵⁷ “Integrity and Impartiality,” 40.

which generates these demands has considered the significance of agency in formulating these demands, we might have no legitimate objection to them. However great, though, these demands will not require any individual to violate her own agency.

The fact that these sacrifices are required is a result of our limitations, of our finite resources and capacities as humans. We simply cannot do everything. Even if we would like to, there comes a point where morality would, without some boundary, ask so much of us that it wouldn't make sense to accede to these demands—acting in accordance with them would require us to draw from resources which are needed to preserve functions central to our own agency. The resources we have, including our attention, time, and energy, are vast and enable us to accomplish a great deal, but these resources are still limited. In many instances, we would be making the best use of our resources by working to structure our norms and institutions such that morality need not require so much of us and so we can do more with our limited resources: our duties ought to include working for social and political arrangements which will encourage ourselves and others to promote the good, and to do so in a way that allows us to make the best use of these limited resources. Even if we do so, however, our limitations will allow us to do much less than we would like.

For some readers, my discussion of agency and its implications for moral requirement may, despite my efforts, seem intolerably vague. If these concepts are to play such a central role in an account of moral duty and its limits, we certainly ought to have a clear picture of what they are and their implications for moral demands. Here I will conclude by trying to trace some more specific and substantive consequences of the recognition of agency for morality's demands and their limits.

Central to the notion of agency is the capacity for rationality. So, morality cannot require us to subvert our own capacity for rational thought. Furthermore, this implies that we should make the development and refinement of our own rational capacities a priority—their effective exercise is necessary to perform our moral duties.

I have also mentioned the importance for agency of forming and maintaining genuine, meaningful attachments. If we do not experience caring, companionship, and even love, we cannot be expected to understand how others might experience these—and thus cannot be expected to make anything like an informed judgment about the moral significance of these concepts in the lives of others. This might seem a cold and moralistic account of the significance of such concepts—indeed, it is—but I am not claiming that this is the only way in which these deep attachments are valuable. I am simply claiming that *one* way in which they are valuable is in their centrality to our functioning as moral agents, and that these relations bear in important ways on the limitations of what we can be morally required to sacrifice. It should certainly be noted that even on the most generous account of the value of such attachments for the functioning of a rational agent, we might not be morally permitted to be a social butterfly—but we might also not be morally permitted to be a hermit crab.

Furthermore, Herman points out that moral demands must reflect the environmental and material factors necessary to cultivate agency:

Our agency arises in ordered stages; it is the result of a process shaped by natural and social resources, completed by our own choices. Within a range of normal variation, there are general conditions for effective agency. Many of the resources that support successful or developing agents cannot be made available for use by others without undermining the agency from which they would be withdrawn....Because effective agency is not like getting one's adult teeth, it will not just happen with time and food, a moral theory that prizes the value of rational agency has

to be especially sensitive to its social and material conditions as it goes about the business of parceling out goods.⁵⁸

A certain amount of resources will be essential to our ability to cultivate and maintain our agency; we need certain basic necessities, as well as resources to pursue the (permissible) projects which form our character—and requirements to sacrifice these resources, which are central to projects necessary for our agency, should be seen as requirements to sacrifice our agency—and, thus, as unacceptable. The activities and resources necessary to pursue these projects will, though, vary greatly from person to person. It is with this aspect of the limitation on our obligations to aid others, however, that I think it most easy to get carried away. We will often be tempted to construe projects as “necessary to develop and sustain agency” when in fact they are not. All I will say here is that we must be exceedingly judicious in our and others’ classifications of activities as central to our agency, and always be willing to accept that we might be morally inclined to give up luxuries to which we have become deeply attached.

Even this attempt to unpack some basic but substantive implications of agency as a limiting condition on our moral obligations may seem too general to be of use. However, I would point out that giving more than a general, qualitative account of the features that are relevant to discussions of moral demandingness is quite difficult, and would probably be futile. Individual situations are so complex, and moral deliberation so subtle, that we will often find that we have to decide for ourselves (or seek help from others in order to determine) when certain sacrifices would put our agency at risk. Indeed, the concept of agency itself is fragmented and multidimensional, and I wish to avoid delving into any exhaustive analysis here, despite the knowledge that my account of its

⁵⁸ “The Scope of Moral Requirement”, 241.

moral significance will be unsatisfying for many. I only wish to offer the reader a schematic sketch in order to give at least some idea of what the basic features of an agency-centered constraint might look like in practice.

I will emphasize a final point relating to moral requirement and our finitude as agents.⁵⁹ A recognition of the cognitive, physical, and economic limitations of our lives must also play an important role in our decisions about what to deliberate about, as we lack the ability to even weigh the balance of reasons with respect to every single action performed. Sometimes, especially when the reasons on both sides are many and weak, it will be quite difficult to tell whether we have legitimate and strong enough reasons not to perform what appears to be our moral obligation. Fair enough, but this need not undermine our faith in our moral deliberation; it need not confuse us about the fact that there *are* important cases in which our desires, our ground projects, or other valued pursuits are clearly dwarfed by the demands of other moral considerations like the immense suffering of others. While we might even accept that in a variety of cases and for an enormous variety of reasons, we will make the wrong decisions, we should take extra care to distinguish cases in which those failures will have major moral ramifications from cases in which our conduct might be of little (but still some) moral consequence. As Goodin puts it, we ought to distinguish the moral “high crimes” from mere “misdemeanors,” and take extra care to ensure that we avoid the former.⁶⁰ As finite agents, we cannot constantly patrol our every action with perfect judgment—our resources, cognitive and otherwise, are simply insufficient for us to do so. We must be sure to worry about the right situations; we can relax our judgments when they are of

⁵⁹ My thinking here stems, in part, from a discussion of this issue in Sarah Buss, “Needs (someone else’s), Projects (my own), and Reasons,” see especially 398-402.

⁶⁰ “Demandingness as a Virtue”, 12.

obviously less significance and we must pay extra attention when we know that, morally speaking, there is a great deal on the line.

Bibliography:

- Arthur, John. "Rights and the Duty to Bring Aid." *World Hunger and Moral Obligation*. Ed. William Aiken and Hugh LaFollette. Englewood Cliffs, N.J.: Prentice Hall, 1977. 37-48.
- Bratman, Michael E. "Kagan on "The Appeal to Cost"." *Ethics* 104.2 (1994): 325-332.
- Brewer, Talbot. Chapter 6: "Against Modern Dualism About The Good", from an untitled forthcoming work (2009).
- Buss, Sarah. "Needs (someone else's), Projects (my own), and Reasons." *The Journal of Philosophy* 103.8 (2006): 373-402.
- Carbonell, Vanessa. *Moral Saints Reconsidered*. Diss. University of Michigan – Ann Arbor, 2009.
- Goodin, Robert. "Demandingness as a Virtue." *Journal of Ethics* 13.1 (2009): 1-13.
- Hare, Caspar. "Rationality and the Distant Needy." *Philosophy and Public Affairs* 35.2 (2007).
- Herman, Barbara. "The Practice of Moral Judgment." *The Journal of Philosophy* 82.8 (1985): 414-436.
- Herman, Barbara. *The Practice of Moral Judgment*. USA: Harvard University Press, 1993.
- Herman, Barbara. "The Scope of Moral Requirement." *Philosophy and Public Affairs* 30.3 (2002): 227-256.
- Hurley, Paul. "The Hidden Consequentialist Assumption." *Analysis* 52.4 (1992): 241-248.
- Jackson, Frank. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101.3 (1991): 461-482.
- Jacobson, Daniel. "Utilitarianism without Consequentialism: The Case of John Stuart Mill." *Philosophical Review* 117.2 (2005): 159-191.
- Kagan, Shelly. *The Limits of Morality*. USA: Oxford University Press, 1989.
- Kagan, Shelly. "Defending Options." *Ethics* 104.2 (1994): 333-351.

- Kagan, Shelly. "Review: Précis of The Limits of Morality; The Limits of Morality." *Philosophy and Phenomenological Research* 51.4 (1991): 897-901.
- Kagan, Shelly. "Review: Replies to My Critics; The Limits of Morality." *Philosophy and Phenomenological Research* 51.4 (1991): 919-928.
- Kagan, Shelly. "Does Consequentialism Demand too Much?" *Philosophy and Public Affairs* 13.3 (1984): 239-254.
- Kleinig, John. "Good Samaritanism." *Philosophy and Public Affairs* 5.4 (1976): 382-407.
- Lichtenberg, Judith. "Negative Duties, Positive Duties, and the "New Harms"." *Ethics*, forthcoming in April 2010.
- Miller, Richard W. "Beneficence, Duty and Distance." *Philosophy and Public Affairs* 32.4 (2004): 357-383.
- Murphy, Liam B. "The Demands of Beneficence." *Philosophy and Public Affairs* 22.4 (1993), 267-292.
- Murphy, Liam B. *Moral Demands in Nonideal Theory*. New York: Oxford University Press, 2000.
- Nagel, Thomas. "Personal Rights and Public Space." *Concealment and Exposure*. Oxford: Oxford University Press, 2002. 31-52.
- Neiman, Susan. *Moral Clarity: A Guide for Grown-Up Idealists*. USA: Harcourt, 2008.
- Nozick, Robert. *Anarchy, State, and Utopia*. USA: Basic Books, Inc., 1974.
- O'Neill, O'Nora. "Kantian Approaches to Some Famine Problems." *Matters of Life and Death: New Introductory Essays in Moral Philosophy*, 3rd Edition. Ed. Tom Regan. New York: McGraw-Hill, 1980. 258-70.
- Portmore, Douglas W. "Commonsense Morality and Not Being Required to Maximize the Overall Good." *Philosophical Studies* 100.2 (2000): 193-213.
- Portmore, Douglas W. "Consequentializing Moral Theories." *Pacific Philosophical Quarterly* 88 (2007): 39-73.
- Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13.2 (1984): 134-171.
- Raz, Joseph. "A Morality Fit for Humans." *Michigan Law Review* 91 (1993): 1297-1314.
- Scanlon, T.M. *What We Owe to Each Other*. Cambridge: Harvard University Press, 1998.

- Scheffler, Samuel. *Human Morality*. New York: Oxford University Press, 1992.
- Scheffler, Samuel. "Morality's Demands and Their Limits." *The Journal of Philosophy* 83.10 (1986): 531-537.
- Scheffler, Samuel. *The Rejection of Consequentialism (Revised Edition)*. USA: Oxford University Press, 1994.
- Schroeder, Mark. "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 117.1 (2007): 265-295.
- Shiffrin, Seana. "Moral Autonomy and Agent-Centred Options." *Analysis* 51.4 (1991): 244-254.
- Singer, Peter. "Famine, Affluence and Morality." *Philosophy and Public Affairs* 1.3 (1972): 229-243.
- Slote, Michael. *Common-sense Morality and Consequentialism*. Great Britain: St. Edmundsbury Press, 1985.
- Slote, Michael. "Shelly Kagan's *The Limits of Morality*." *Philosophy and Phenomenological Research* 51.4 (1991): 915-917.
- Smart, J.J.C. and Bernard Williams. *Utilitarianism for & against*. USA: Cambridge University Press, 1973.
- Unger, Peter. *Living High and Letting Die*. New York: Oxford University Press, 1996.
- Urmson, J.O. "Saints and Heroes." *Essays in Moral Philosophy*. Ed. Abraham Irving Melden. USA: The University of Washington Press, 1958. 198-216.
- Waldron, Jeremy. "A Right to Do Wrong." *Ethics* 92.1 (1981): 21-39.
- Watson, Gary. "Skepticism about Weakness of Will." *The Philosophical Review* 86.3 (1977): 316-339.
- Williams, Bernard. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press, 1985.
- Williams, Bernard. *Moral Luck*. New York: Cambridge University Press, 1981.
- Wolf, Susan. "Above and Below the Line of Duty." *Philosophical Topics* 14.2 (1986): 131-148.
- Wolf, Susan. "Moral Saints." *The Journal of Philosophy* 79.8 (1982): 419-439.