

MAKING ROOM FOR LOVE:
A MODEST DEFENSE OF IMPARTIAL MORALITY

Keshav Singh

A thesis submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Arts in the Department of Philosophy.

Chapel Hill
2016

Approved by:

Susan Wolf

Russ Shafer-Landau

Geoff Sayre-McCord

© 2016
Keshav Singh
ALL RIGHTS RESERVED

ABSTRACT

Keshav Singh: Making Room for Love: A Modest Defense of Impartial Morality
(Under the direction of Susan Wolf)

There is a *prima facie* tension between impartial morality, on one hand, and the kind of partiality necessary for love, on the other. In this paper, I explore this tension in light of critiques of impartial morality by Michael Stocker and Bernard Williams, and argue that their remarks show it to be a deep and difficult tension. I go on to examine two kinds of attempts to resolve this tension, and argue that they fail to reconcile impartial morality with the kind of partiality necessary for love. I then offer my own solution to the problem, according to which impartial morality can be reconciled with the right kind of partiality as long as we grant that morality is not always overriding. When we grant this, I argue, we make room for love, because we are not forced to have standing unconditional commitments to acting in accordance with impartial moral principles.

ACKNOWLEDGEMENTS

First, I would like to thank my adviser, Susan Wolf. Her importance to this project is difficult to overstate. She devoted a great deal of time to discussing my ideas with me and giving me feedback on my writing. Moreover, the influence of her philosophical ideas on this paper does not even come close to being captured by my citations of her work. Her writing has set the agenda for talking about love in philosophy, and is one of the primary inspirations for this project.

I would also like to thank Russ Shafer-Landau and Geoff Sayre-McCord for serving on my committee. Their comments and suggestions were essential to this paper becoming what it is now.

I would like to thank my parents and brother for their support and encouragement. All of them are paragons of scholarship and intellectual achievement, and I have achieved what I have academically by following their examples.

Lastly, I would like to thank Lindsay Brainard. Her patience, understanding and unwavering support helped me through my moments of difficulty, and her willingness to discuss my ideas and proofread my work contributed greatly to the quality of this paper.

TABLE OF CONTENTS

INTRODUCTION.....	1
SECTION I: THE PROBLEM WITH IMPARTIALITY.....	2
SECTION II: TWO-LEVEL THEORIES.....	7
Two-level Consequentialism.....	9
Two-level Kantianism.....	12
SECTION III: “SOPHISTICATED” THEORIES.....	15
Sophisticated Consequentialist Partiality.....	16
Sophisticated Kantian Partiality.....	20
SECTION IV: IMPARTIALITY AND THE AUTHORITY OF MORALITY.....	24
Overridingness and All-Things-Considered.....	25
A Modest Defense of Impartiality.....	28
The Ordering of Practical Perspectives.....	34
CONCLUSION.....	38
REFERENCES.....	39

Introduction

There is a *prima facie* tension between two theses, each of which is not only highly intuitively plausible on its own, but held dear by many moral theorists. On one hand, there is the thesis that morality is impartial, which is a part of both the consequentialist and Kantian approaches that have dominated moral philosophy for the last few hundred years. On the other hand, there is the thesis that partial action is sometimes permissible, or even required. The costs of rejecting the latter thesis are extremely high, because partial action plays such an indispensable role in the formation and cultivation of the kinds of interpersonal relationships that, for many of us, make life worth living. In other words, partial action plays an indispensable role in *love*, broadly construed. Thus, it may seem that, if these two theses are truly incompatible, we must reject the thesis that morality is impartial. But that thesis is also not easy to reject, for it may also seem like the kind of equal regard for all persons constitutive of impartiality is also partly constitutive of morality. Thus, if there is a solution in the offing that shows these two theses to be compatible, that would, I take it, be preferable to rejecting the impartiality of morality.

In this paper, I will attempt to offer such a solution. In the first section, I will flesh out the tension between partiality and impartial morality, and argue that critiques of consequentialism and Kantianism presented by Michael Stocker and Bernard Williams show the difficulty of resolving this tension. In the second and third sections, I will discuss two sorts of attempts by consequentialists and Kantians to address this issue. The first set of attempts consists of

arguments that partial action can be justified on ultimately impartial grounds. The second set of attempts consists of arguments that impartial moral requirements serve as limiting conditions on our actions, but are not the only sources of justification for them. I will argue that both sorts of attempts are unsuccessful. In the fourth section, I will attempt a different sort of resolution, which I call a “modest defense” of impartial morality. According to this modest defense, impartial morality can coexist with the right kind of partiality as long as we concede that, from the perspective of what we ought to do all-things-considered, moral requirements are not always overriding.

I. The Problem with Impartiality

Many of us have the conviction that morality is an *impartial* system. There are many ways of spelling out exactly what it amounts to for morality to be impartial, but I am not concerned with those specifics in this paper. For the time being, all that is necessary is a general idea of what impartiality is that (a) is more or less neutral between various more specific accounts of impartiality, and (b) gets to the heart of the conviction that morality is impartial. I think the following satisfies these two *desiderata*: the thesis that morality is an impartial system is the thesis that from a moral perspective, every moral agent has equal standing, and thus, at least at the most fundamental level, each “counts” just as much as any other.¹ That morality has this feature is, I think, a very attractive view, because it makes sense of our intuition that morality is a fair system that treats everyone equally.

¹ There may be other members of the “moral community” that are not moral agents, like non-human animals, but I cannot address here the question as to what moral standing animals have.

However, most of us also have the conviction that it's at least sometimes permissible, or perhaps even required for people to act partially – that is, to give priority to certain people over others in their actions. On the face of it, these two convictions seem to be in conflict; if morality requires equal regard for all persons, how can it be permissible to give priority to some people over others? The most obvious explanation we can give of why we are justified in giving priority to some people over others is that those people count more than others. But if morality is impartial, it seems this explanation must be false, or at least too simple. To theorists who are already familiar with various accounts of how impartial morality can accommodate partiality, the tension between partiality and impartial morality may be thought to be not particularly deep or difficult to resolve. But pre-theoretically, at least, the tension does exist, and it is precisely the kind of tension that defenders of impartial morality are often worried about.

I think this tension is shown to be a deep and difficult one by Michael Stocker and Bernard Williams, among others, in their critiques of consequentialism and Kantianism². Though Stocker and Williams each present distinct critiques of these theories, I think there is a point of great importance to be taken from both of them, that at the core of our convictions about the justification of partiality is a conviction about *love*, broadly construed. Partial action, in general, seems to be a necessary part of a variety of interpersonal relationships. But intuitively, there is a difference not just in degree but in kind between the partiality that can be said to be exhibited by a *quid pro quo* exchange, and that which figures into the deepest, closest relationships, whether they be with friends, relatives, or romantic partners. Let us call this latter kind of relationship an *authentic loving relationship*, because it is the kind of relationship that is not just apparently but

² See Stocker (1976) and Williams (1973, 1981).

actually based on love. Part of the reason we are so concerned about partiality is that there is a certain kind of partiality that is partly constitutive of these authentic loving relationships.³ In other words, part of why we want to make room for partiality in our moral theories is that we want to make room for love. Thus, it is not enough for defenders of morality to show that partiality of *some kind* is compatible with impartial morality. It must be shown that partiality of the right kind – the kind that is partly constitutive of authentic loving relationships – is compatible with impartial morality. As I will argue, this is a deep problem for both Kantianism and consequentialism *qua* impartial moral theory.

Stocker and Williams each put the point differently, so I think it is worth briefly discussing what each of them says. According to Stocker, there is a dilemma that arises for both consequentialists and Kantians: either one must act on the basis of the reasons provided by one's moral theory, or one must act on the basis of something else.⁴ If one acts on the basis of the reasons provided by consequentialism or Kantianism, one is unable to have authentic loving relationships. If she does not act on the basis of these reasons, she is “morally schizophrenic,” because there is a split between her reasons and her motives – in other words, she is not motivated by what she thinks she has reason to do. The important point here is that, in Stocker's view, actions performed for Kantian or consequentialist reasons are not capable of cultivating

³ As I will discuss later, I take partiality to be a necessary part not only of authentic loving relationships, but also of an important kind of self-love.

⁴ When I say “on the basis of,” I mean to refer to something like the basing relation discussed in epistemology. The epistemic basing relation holds between a belief and a reason just in case that is a reason *for which* the belief is held. An analogous basing relation can be said to hold in the practical case, except the relation holds between an action (or perhaps an intention, but I don't think it matters for the purposes of this paper) and a reason just in case that is a reason *for which* the action is performed.

authentic loving relationships. When one acts so as to fulfill one's duty, or to maximize the good, one is not acting out of love

...for it is essential to the very concept of love that one care for the beloved, that one be prepared to act for the sake of the beloved. More strongly, one must care for the beloved and act for that person's sake as a final goal; the beloved, or the beloved's welfare or interest, must be a final goal of one's concern and action.⁵

Thus, Stocker contends, there is a tension between acting for Kantian or consequentialist reasons, and acting out of love. And insofar as acting out of love requires acting partially, there is a worry that even if one is able to act partially for Kantian or consequentialist reasons, this will not be good enough to solve the problem.

Williams makes similar points in several places in his writing, sometimes putting the issue in terms of integrity or alienation, among other things. Probably the most famous (and most poignant) of these arguments is that justifying one's partial actions by appeal to general moral principles provides the agent with "one thought too many."⁶ Williams gives the example of a man saving his drowning wife instead of a stranger, pointing out that if he justifies saving her on the grounds that doing so is legitimated by some general moral principle, there seems to be something missing in the way he relates to her. She might have hoped that the fact that it was *her* drowning would be enough for her husband to choose to save her. Thus, the thought that saving her is legitimated by some general moral principle is one thought too many for the man to have. Importantly for my purposes, Williams ties his thought experiment explicitly back to impartial morality:

⁵ Stocker, p. 456.

⁶ Williams (1981), p. 18.

Perhaps others will have other feelings about this case. But the point is that somewhere (and if not in this case, where?) 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.⁷

Though he is somewhat reticent about the details, Williams' point here seems to be that there is a deep conflict between acting partially in a way that expresses one's deep commitment to another person and acting from the perspective of impartial morality. Like Stocker, I think Williams shows that impartial morality must be reconciled not with partial action *simpliciter*, but with the right kind of partiality.

Both Stocker and Williams criticize consequentialism and Kantianism as being at odds with the kind of partiality that is constitutive of authentic loving relationships. Though Stocker never puts his critique explicitly in terms of impartiality like Williams does, I take impartiality to be the objectionable feature common to the two theories they discuss. Both Kantianism and consequentialism are impartial moral theories in the sense I described earlier, and it is natural to think of Stocker's and Williams' critiques of the two theories partly as pointing to deeper, more difficult versions of the *prima facie* tension between impartial morality and the important role that partial action plays in our lives. What I will take from Stocker and Williams, then, is that if we are to properly resolve this tension, we must resolve it in a way that shows impartial morality to be compatible with the kind of partiality that is partly constitutive of authentic loving relationships. With this in mind, I will spend the next two sections critiquing two sets of responses to the tension by defenders of Kantianism and consequentialism, respectively.

⁷ Williams, *op. cit.*

II. Two-Level Theories

One way in which defenders of consequentialism and Kantianism try to respond to such critiques is to attempt to accommodate partiality within their theories. They do so by arguing that partial actions can be morally justified on the basis of impartial consequentialist or Kantian moral principles. Thus, theories that are fundamentally impartial can accommodate partiality at the level of particular actions or even kinds of actions. When writing about the objection that impartial moral theories can't allow partiality, Marcia Baron elegantly summarizes the appeal of such responses:

But it is, I think, rooted more deeply in an error, specifically, in a failure to distinguish different points or "levels" at which impartiality might be deemed requisite. Critics suppose that impartialists insisting on impartiality at the level of rules or principles are committed to insisting on impartiality at the level of deciding what to do in one's day-to-day activities. (The critics of impartialists are helped into this error by some of the extreme impartialists themselves, e.g., Peter Singer.)⁸

The "extreme impartialists" Baron refers to insist on impartiality at the level of actions (or kinds of actions) and at the level of general principles. It seems to me that such theorists, because they endorse theories that demand impartiality at both levels, have no response to the objection but to bite the very large bullet. On the other hand, theorists whom we might term "moderate impartialists" insist on impartiality only at the level of general principles, not at the level of actions or kinds of actions. On the moderate impartialist view, while fundamental moral principles must all be such that they give all moral agents equal standing, these principles can give us moral reasons to give priority to some people over others. Because on such views partial

⁸ Baron (1991), p. 842.

actions can be morally justified, moderate impartialists may seem to be in a much better position to respond to the critiques in question.

However, I think their responses still fall short. Baron is right to point out that critics of impartial morality often fail to distinguish between two levels at which partiality might be accommodated. Being able to accommodate partiality at the level of actions is certainly better than not being able to accommodate it at all. However, I think we must also distinguish between two *kinds* of justification for partial action. The first kind justifies partial action on the basis of impartial general principles; because the justification for such actions is not itself partial, I will call the partial status of such actions *surface partiality*. The second kind does not justify partial action on the basis of anything impartial; in fact, it may not justify it on the basis of any general principles at all. But whatever the basis for the justification is, it is itself partial. I will refer to the partial status of such actions as *deep* partiality. I contend that accommodating surface partiality is not enough to successfully respond to Stocker and Williams' critiques, because it is only through deep partiality that we can relate to friends and loved ones in the right ways. In this section, I will consider both a consequentialist attempt and a Kantian attempt to accommodate partiality within their theories. I will argue that while both attempts successfully accommodate surface partiality, the fact that they cannot accommodate deep partiality prevents them from responding successfully to the critiques in question.

Two-Level Consequentialism

Consequentialists often argue that we have moral reasons to perform many partial actions on the grounds that they actually end up maximizing the good.⁹ For example, utilitarians have long argued that since people are best at maximizing the happiness of those closest to them, the overall good is maximized when everyone acts partially toward their friends and loved ones. This is the kind of argument Sidgwick offers for the moral justification of partial action:

And certainly all special affections tend occasionally to come into conflict with the principle of promoting the general happiness; and Utilitarianism must therefore prescribe such a culture of the feelings as will, so far as possible, counteract this tendency. But it seems that most persons are only capable of strong affections towards a few human beings in certain close relations, especially the domestic : and that if these were suppressed, what they would feel towards their fellow-creatures generally would be, as Aristotle says, “ but a watery kindness “ and a very feeble counterpoise to self-love : so that such specialised affections as the present organisation of society normally produces afford the best means of developing in most persons a more extended benevolence, to the degree to which they are capable of feeling it. Besides, each person is for the most part, from limitation either of power or knowledge, not in a position to do much good to more than a very small number of persons; it therefore seems, on this ground alone, desirable that his chief benevolent impulses should be correspondingly limited.¹⁰

According to Sidgwick, there are two reasons why happiness is maximized when people act partially toward their friends and loved ones. The first is that the “strong affections” that motivate people toward the most benevolent actions can only be felt by most people toward a small group of friends and loved ones, so fostering these strong affections seems to be the best way of ensuring that one will act benevolently, even if it means that one will do so disproportionately toward a small group of people. The second is that due to their relationships

⁹ I will only discuss act-consequentialism in this paper, but as far as I can tell, all of the arguments apply to rule-consequentialism as well.

¹⁰ Sidgwick, p. 434.

with friends and loved ones, most people are simply in a much better position to do good for those people. Thus, according to utilitarianism, we have moral reasons for continuing to maintain close relationships with friends and loved ones and act partially toward them.

The character John, from Peter Railton's "Alienation, Consequentialism and Morality" gives us a sense of what it would look like for someone to embrace this utilitarian justification of partiality:

To many, John has always seemed a model husband. He almost invariably shows great sensitivity to his wife's needs, and he willingly goes out of his way to meet them. He plainly feels great affection for her. When a friend remarks upon the extraordinary quality of John's concern for his wife, John responds without any self-indulgence or self-congratulation, "I've always thought that people should help each other when they're in a specially good position to do so. I know Anne better than anyone else does, so I know better what she wants and needs. Besides, I have such affection for her that it's no great burden – instead, I get a lot of satisfaction out of it. Just think how awful marriage would be, or life itself, if people didn't take special care of the ones they love." His friend accuses John of being unduly modest, but John's manner convinces him that he is telling the truth: this is really how he feels.¹¹

As Railton notes, there is something troubling about John's explanation of why he acts partially toward Anne. We would expect out of an authentic loving relationship that John acts, as Stocker would put it, *for the sake* of Anne; instead, he acts partially toward her only because he realizes that it's the best way for him to contribute to the overall good. In other words, in acting partially toward her for consequentialist reasons, he acts with mere surface partiality toward her. And it seems like this surface partiality is not enough for John to be able to have an authentic loving relationship with Anne.

To tie this example back to Williams' point about "one thought too many," Anne might expect that if John really loved her, considerations like "it's my wife" or "it's *Anne*" would play

¹¹ Railton, p. 135.

an unmediated role in his explanation of why he acts partially toward her. John's explanation of his partial actions toward Anne does reference the fact that it's *her*, but it is not just on the basis of this kind of consideration that John acts. Instead, the role this kind of consideration plays is mediated by an impartial general principle (the utilitarian criterion of rightness). It is not just because it's Anne that John acts partially toward her, but also because he is in a special position to satisfy her wants and needs that puts him in a position to maximize the good if he does so. If utilitarianism is the correct moral theory, then the impartial utilitarian criterion of rightness is the only possible source for any moral reason he has to act partially toward Anne. Thus, if John is a wholehearted utilitarian, he must believe that this principle is what morally justifies his partial actions toward Anne (whether or not he uses this principle as a decision procedure).

It seems quite clear that this kind of attempt by the consequentialist to accommodate partial action does not successfully address the problem raised by Stocker and Williams. While it may show that there can be consequentialist reasons for partial action, these are not the right kinds of reasons because they are grounded in impartial general principles. In acting partially toward another for consequentialist reasons, one acts for a reason that gets its normative force not from facts like "it's her", but from the fact that doing so will maximize the good. This is a problem for all such consequentialist justifications; nothing about it turns on whether one is a utilitarian or a value-pluralist. For even if a consequentialist were to take acting partially toward loved ones to be intrinsically good, she would still be acting partially toward them only insofar as doing so maximized the good, and thus acting only with surface partiality. Thus, it seems that the two-level consequentialist approach, at least, cannot accommodate the right kind of partiality.

Two-level Kantianism

A Kantian might offer a similar kind of “moderate impartialist” response to Stocker’s Williams’ critiques. According to Kantianism, an action is morally permissible if and only if it is in accordance with the categorical imperative, which is, of course, an impartial principle. But the categorical imperative is meant to do more than just offer constraints on what is morally permissible; it generates not just negative duties, but positive duties as well. Thus, when considering whether she has a moral reason to act partially toward another person, the Kantian must consider not only whether doing so would be *permissible*, but also whether the categorical imperative generates a positive duty for her to do so. For example, when considering the moral status of parental partiality, the Kantian must consider not only whether a maxim enjoining each parent to take care of her own children is universalizable, but whether treating her children as ends in themselves generates a duty for her to take care of them.

In the *Metaphysics of Morals*, Kant attempts to draw out some of the positive duties generated by the categorical imperative, including the duty to take care of one’s children. Kant argues that some degree of parental partiality is not just morally permissible, but morally required:

So from a *practical point of view*, it is a quite correct and even necessary idea to regard the act of procreation as one by which we have brought a person into this world without his consent and on our own initiative, for which deed the parents incur an obligation to make the child content with his condition so far as they can. They cannot destroy their child as if he were something they had *made* (since a being endowed with freedom cannot be a product of this kind) or as if he were their property, nor can they even just abandon him to chance, since they have brought not merely a worldly being but a citizen of the world into a condition which cannot now be indifferent to them even just according to concepts of right.¹²

¹² Kant, p. 64.

According to Kant, parents have moral obligations to take care of their own children that they don't have toward other children, which they incur by bringing their children into the world. However, though this is an obligation to act partially toward one's children, it doesn't seem to be an obligation to do so for their own sake. Rather, it seems like for Kant, the normative force of this obligation comes from a more general, impartial principle about what we owe to those we bring into this world without their consent. Just as there can be impartial principles that dictate what we owe to those to whom we have incurred debts, this too looks to be an impartial principle. And all such principles, on Kant's view, ultimately fall out of the fundamental impartial principle of the categorical imperative.

As Thomas Hill writes, on Kant's view, "the fact that a principle would benefit *me*, *my* family, and *my* country instead of someone else, someone else's family and country, is not *in itself* a reason for anyone, as a moral legislator, to favor that principle. In other words, at the level of deliberation about basic principles, morality requires impartial regard for all persons."¹³ This supports the interpretation of Kant's remarks on parental partiality and other kinds of partial action that they only accommodate partiality through appeal to ultimately impartial principles. Hill, making a similar point to Baron's, goes on to write the following:

Autonomy, as impartiality, is part of an ideal for moral legislation, or general debate about moral principles and values; it is not a recommended way of life. Unfortunately, some philosophers, including Kant, seem at times to conflate this legislative ideal with another idea, which does not really follow: namely, the idea that in facing the moral choices of daily life we should constantly strive to act on impartial principles, to free ourselves of particular attachments, and to ignore the distinguishing features of individuals.¹⁴

¹³ Hill, p. 45.

¹⁴ Hill, op. cit.

Hill's remarks suggest that Kantianism should be formulated as a moderately impartial theory, on which though the theory is impartial at the level of fundamental moral principles, it need not be impartial at the level of actions or even kinds of actions. Returning to the example of parental partiality in Kant's ethics, it seems, therefore, that though partial actions are often morally permissible, and even required, their status as such is grounded in impartial moral principles. I will argue that because of this, as with the two-level consequentialist response, this two-level Kantian response justifies only surface partiality.

Consider a parent, Asha, who, following this kind of Kantianism, acts partially toward her son Jay in various ways. She shows great sensitivity to Jay's wants and needs, and often goes out of her way to fulfill them. Furthermore, she invests a great deal of resources in his development that she does not invest in the development of the children of others. When Asha's friends commend her on what a model parent she is to Jay, she responds, somewhat self-effacingly, "Of course I care about Jay, but I'm just doing my duty. Jay didn't decide to be born – it was my decision to bring him into this world. So I owe it to him to give him a good life and make sure he develops into a good person. I'm just doing what any parent is morally obligated to do for her children." The example of Asha and Jay should, I think, remind us of Railton's example of John and Anne; here too, there is something missing in the way Asha relates to Jay. While she acts partially toward Jay in a variety of ways and those actions are, according to Kantianism, not only morally permissible but morally obligatory, she acts out of a sense of duty founded on impartial moral principles about what parents owe to their children. In other words, she acts with surface partiality, not deep partiality.

Again, we might expect from a parent who truly loves her child that the kinds of considerations Williams discusses, in the form of “it’s my *son*” or “it’s *Jay*,” would play an unmediated role in Asha’s explanation of why she acts partially toward him. However, though the fact that Jay is her son plays an essential role in her explanation of why she acts partially toward him, that role is mediated by an impartial general principle, because it is only insofar as the categorical imperative generates a duty for her to do so that she thinks she has such a duty. It may be that Asha feels a great deal of affection toward Jay; however, according to her Kantian beliefs, this fact plays no role in the moral justification of her partiality toward Jay. It seems what is missing in this parent-child relationship is *deep* partiality – partial action that is itself justified on partial grounds. This example is meant to illustrate what is deficient about the way in which any Kantian, who acts partially toward loved ones out of a sense of her ultimately impartial moral duty, relates to those loved ones. Thus, because like the two-level consequentialist response, the two-level Kantian response can accommodate only surface partiality, it does not succeed in addressing Stocker’s and Williams’ critiques.

III. “Sophisticated” Theories

In the previous section, I considered attempts to accommodate partiality within impartial moral theories by showing that partial action can be morally justified on the basis of impartial moral principles. I argued that these attempts fail because they can accommodate only what I have called surface partiality within their theories, not deep partiality. Thus, they fail to respond to the charges from Stocker and Williams that their theories prohibit authentic loving relationships. However, there is another, more sophisticated way (as Railton would put it) in

which defenders of impartial morality, both consequentialist and Kantian, attempt to accommodate partiality within their theories. Such attempts begin with the claim that there is nothing in their theories that requires one always to act on the basis of impartial moral principles. Certainly, the moral status of one's actions is determined by impartial moral principles, and impartial moral theories require a commitment on the part of the agent that her actions accord with those principles. But it does not follow from these features of impartial moral theories that one is prohibited from acting on bases (such as love) other than these principles. Thus, impartial moral theories can accommodate deep partiality. In this section, I'll consider two responses of this nature, one consequentialist and one Kantian. I'll argue that these responses fail, not because they fail to accommodate deep partiality, but because they require an unconditional commitment to acting in accordance with impartial principles that renders their partiality toward loved ones too fragile to constitute authentic loving relationships.

Sophisticated Consequentialist Partiality

Earlier, I discussed an attempt to accommodate partial action within consequentialist theories that utilitarians like Sidgwick have argued for. I argued that this attempt fails, because in justifying partial action on the basis of an impartial moral principle (the utilitarian criterion of rightness), one justifies only what I have called surface partiality. As I have argued, the right kind of partiality must be deep partiality. Railton makes a similar point, using the example of John and Anne, which I discussed earlier, to show that there is something missing in the way John relates to Anne. Thus, presumably, Railton agrees that the Sidgwick-style response fails to accommodate partiality in a way that adequately addresses Stocker's and Williams' worries.

Nevertheless, he thinks that consequentialism has the resources to respond to such worries. He suggests that consequentialism does not require that its adherents always follow a consequentialist decision procedure; rather, a *sophisticated* consequentialist is committed to following whatever decision procedure leads her to live an objectively consequentialist life (a life in which one maximizes the good), even if that means acting on the basis of something other than the consequentialist criterion of rightness. Thus, Railton argues, the sophisticated consequentialist can, without being at odds with consequentialism, act out of love, as long as doing so is compatible with her commitment to living an objectively consequentialist life.

To draw this point out, Railton has us consider the example of another couple, Juan and Linda. Juan, like John, has always seemed to be a model husband. However, he offers a different explanation of his actions from John's:

Look, it's a better world when people can have a relationship like ours and nobody could if everyone were always asking themselves who's got the most need. It's not easy to make things work in this world, and one of the best things that happens to people is to have a close relationship like ours. You'd make things worse in a hurry if you broke up those close relationships for the sake of some higher goal. Anyhow, I know that you can't always put family first. The world isn't such a wonderful place that it's OK just to retreat into your own little circle.¹⁵

According to Railton, because Juan is a sophisticated consequentialist, he is able to act partially toward Linda *for her own sake*, as long as doing so is constrained by an overarching commitment to leading an objectively consequentialist life. In other words, while Juan is committed to living a life that accords with an impartial general principle (the consequentialist criterion of rightness), he need not always act on the basis of this principle. It seems, therefore, that his partial actions

¹⁵ Railton, p. 150.

toward Linda do admit of deep partiality. If this is the case, then this kind of consequentialism, despite being an impartial moral theory, is compatible with deep partiality.

However, I think there is still something troubling about the way Juan relates to Linda. As Railton notes, Juan's commitment to living an objectively consequentialist life consists partially in his acceptance of the following counterfactual condition: "while he ordinarily does not do what he does simply for the sake of doing what's right, he would seek to lead a different sort of life if he did not think his were morally defensible."¹⁶ Thus, if Juan were to find out that continuing to act partially toward Linda were incompatible with maximizing the good, he would seek to live a life in which he did not act partially toward her, effectively ending their relationship. His partiality toward Linda, then, is in an important sense highly *fragile*, for its continuation depends on the contingent fact that it is compatible with maximizing the good. For example, imagine that Juan learns that he is an incredibly talented matchmaker; he has an almost magical ability to make matches that result in deeply fulfilling relationships like the one he has with Linda. He knows that people having relationships like this makes the world a better place. So if he can sacrifice his relationship with Linda to spend all his time matchmaking and ensure that as many other people as possible can have such valuable relationships, then he is committed to doing so, for staying with Linda would be incompatible with leading an objectively consequentialist life.

We can consider another such example: imagine that Linda develops some horrible illness that renders her so incapacitated that Juan's relationship with her now involves devoting most of his time to serving as her caretaker. Moreover, the nature of her illness is such that taking care of her is extremely resource intensive. In this case, it seems quite unlikely that Juan's

¹⁶ Railton, p. 151.

staying with her is compatible with maximizing the good. Juan would do a much better job of maximizing the good if he instead spent that time in a clinic in a developing country, perhaps using the same amount of resources to save hundreds of people from debilitating diseases. Here, he would also be committed to ending his relationship with Linda (assuming that this were psychologically possible for him). Again, Juan's unconditional commitment to living a life that accords with the consequentialist criterion of rightness comes with a correspondingly fragile commitment to Linda that seems to cheapen their relationship in an important way. The point here is not that authentic loving relationships require unconditional commitments to loved ones, or even that Juan necessarily ought, all things considered, to stay with Linda if she develops a horrible illness. Rather, the point is that there is a problem with consequentialism insofar as it does not even leave Juan with the *option* of staying with his supposedly beloved wife if he discovers that doing so is no longer compatible with the impartial demand of maximizing the good.

Admittedly, Juan's relationship with Linda does not seem as deficient as John's relationship with Anne, because Juan acts with deep partiality toward Linda. However, as I've attempted to show, it is nevertheless deficient in an important way. If he is a wholehearted sophisticated consequentialist, Juan's standing commitment to living in accordance with the impartial demands of consequentialism is such that his relationship with Linda, or any other loved one, will take a back seat to those demands any time he becomes aware of a conflict between them. This is, I think, an example of a problem that will arise for any sophisticated consequentialist, whether or not her relationships with loved ones in fact do come into conflict with the demands of consequentialism. For the problem is not just that in cases where they do come into conflict, the sophisticated consequentialist must give up her relationships. At its heart,

rather, the problem is that the sophisticated consequentialist, in virtue of her unconditional commitment to the impartial demands of consequentialism, is antecedently committed to choosing those demands over her relationship any time they come into conflict, whether or not they in fact do. This, I think, would be an appropriate source of dismay for her loved ones, to find out that their relationship would not persist if not for its lucky coincidence with these demands. We want, and not unreasonably so, more robust commitments from our authentic loving relationships than Railton's consequentialism allows for. Thus, taking the sophisticated consequentialist route still does not allow for the kind of partiality necessary to adequately address Stocker's and Williams' worries.

Sophisticated Kantian Partiality

A similar response to Railton's can be given on behalf of Kantianism. Earlier in this paper, I discussed a form of Kantianism on which it seems that deep partiality is not possible, because in acting partially toward loved ones, one must still act on the basis of impartial principles. However, many self-described Kantians dispute the claim that Kantian theories must include such a requirement. While Kantianism holds that one must act in accordance with the categorical imperative, this is distinct from holding that whenever an agent performs some action, she must act *on the basis* of the categorical imperative. In other words, it isn't the case that the Kantian must offer, as the ultimate justification of each of her actions, an explanation of how that action accords with the categorical imperative. Rather, she must have a commitment to acting in accordance with the categorical imperative that will serve as a limiting condition on how she acts. Thus, she can, without being at odds with her Kantianism, act out of love, as long

as doing so is compatible with her commitment to acting in accordance with the categorical imperative.

Marcia Baron, responding to Stocker and Williams, defends such a view about acting from duty:

On my view, one should (ideally) always act from duty, but this is only to say that all of one's conduct should be governed by one's unconditional commitment to doing what one morally ought to do. To say that one should always act from duty is not to say that one should always act from duty as a primary motive. One's sense of duty will serve generally as a limiting condition and at the same time as an impetus to think about one's conduct, to appraise one's goals, to be conscious of oneself as a self-determining being, and sometimes to give one the strength one needs to do what one sees one really should do.¹⁷

Baron's point here is that while Kantianism requires one always to do one's duty, it does not require one to do one's duty, at the moment of action, *because* it is one's duty. It is not at odds with being a Kantian to act on the basis of love, as long as one's doing so is constrained by an overarching commitment to doing what one morally ought to do. Thus, it seems this kind of Kantianism is compatible with deep partiality, despite being an impartial moral theory.

To draw this out, we can look at an example similar to Railton's example of Juan and Linda. Consider a variation of the case of Asha and Jay discussed earlier in this paper. Now, however, let's imagine that Asha offers a different explanation of her partial actions toward her son. When commended on how attentive she is to Jay's wants and needs, she responds, "Well, of course I am. He's my son and I love him, so I want to give him a good life and make sure he develops into a good person. I would want any mother to do the same for her child." Her response suggests that she acts partially to Jay on the basis of her love for him. Moreover, there seems to be nothing in the kind of Kantianism described by Baron that is at odds with her doing

¹⁷ Baron (1984), p. 209.

so. As long as acting out of love does not conflict with her commitment to doing her moral duty, there seems to be no barrier posed by Kantianism to Asha's loving her son, and acting with deep partiality toward him. But of course, we must remember that this does not mean that Asha's love for her son is unmediated by her commitment to Kantianism. If someone were to ask Asha, "I see why you do so much to care for Jay, but why is it morally permissible for you to do so?" she would have to respond with something like "because it's compatible with doing my duty," or "because it's permitted by the categorical imperative." And then, the question looms: "what if things changed such that the two were no longer compatible?"

It seems clear that if things changed such that acting partially toward Jay were no longer compatible with doing her moral duty, Asha, insofar as she is a wholehearted Kantian, would choose to do her duty over acting partially toward Jay, because her commitment to doing her duty serves as a limiting condition on her actions. In any case where duty required Asha to do something not in Jay's best interest, or even to abandon him, she has committed herself to doing so (antecedently to any such situation arising). Thus, although her partiality toward Jay isn't mediated in the sense that she acts partially toward him ultimately for the sake of duty, it is mediated in the sense that although she acts partially toward him for his sake, she does so only insofar as duty allows. As Baron says, "part of what it is to act from duty is to act with a counterfactual condition always at hand (though not always in one's thoughts): one would not do this if it were morally counter-recommended."¹⁸ Thus, although Kantianism, as an impartial moral theory, can still accommodate deep partiality, it does so with the important caveat that that partiality must be in a sense *fragile* – that is, that the partial actions underlying one's

¹⁸ Baron (1984), pp. 216-217.

commitments to loved ones must not persist in the face of conflict with what duty demands. And this is something that also may seem to be a barrier to having authentic loving relationships – that the agent will have a standing commitment to choose the demands of duty over her relationship if those demands ever come into conflict.

Baron, however, does not see the comparative fragility of the partiality allowed by Kantianism as a problem:

If opponents insist that such a counterfactual condition gets in the way of friendship and love, then, I want to say, so much the worse for friendship and love-thus conceived. Unless one buys into an extremely romantic notion of love and a similar notion of unconditional friendship, the conflict is bogus.¹⁹

Baron thinks that it is no threat to authentic loving relationships that they are subject to the counterfactual condition that one would not act partially toward loved ones if doing so were at odds with the demands of duty. For surely an unconditional commitment to acting partially toward a loved one, no matter what duties would be neglected in the course of doing so, would be bizarre. Thus, such an unconditional commitment is too strong to be a necessary condition for an authentic loving relationship. However, this is a false dilemma; the options here are not restricted to Baron's view and a view on which love requires unconditional commitment to the beloved. Williams himself dismisses as "absurd" and "pathological" the putative commitment to another on which one "must prefer any possible demand of that [relationship] over other, impartial moral demands."²⁰ The problem with Baron's view is not that it places any limits at all on what is permissible to do for the ones we love; rather, the problem is that on her view, we

¹⁹ Baron (1984), p. 217.

²⁰ Williams (1981), p. 17.

must, antecedent to any actual conflicts, have a standing commitment to choose the demands of duty over our relationships with loved ones *whenever* they conflict.

Thus, the objection to Baron's Kantian response is not that it prevents us from having authentic loving relationships if what we do for loved ones is constrained in any way by impartial moral principles. Rather, it is that the constraint, in requiring a standing unconditional commitment to the demands of duty, makes our relationships too fragile in comparison. I concede that a fragile relationship is not as deficient as the kind of relationship discussed earlier that only involves surface partiality. But there is still something less than ideal about the way in which one relates to a loved one if one's commitment to her would, in each and every case of conflict, take a back seat to the demands of impartial morality. It seems to me a deeply disturbing thought that, at best, our loving relationships continue due to their lucky coincidence with these demands. In order for Kantianism (or consequentialism) to allow for the kind of partiality necessary for authentic loving relationships, it must allow for partiality that is not just deep, but appropriately robust.

IV. Impartiality and the Authority of Morality

In the previous sections, I discussed two sets of attempts to reconcile partiality with impartial moral theories, and argued that both fail to allow for the right kind of partiality. The first set, which argues that partial action can be justified on the basis of impartial moral principles, fails because it accommodates only what I have called surface partiality. The second set, which argues that impartial moral principles are not the only sources of justification for partial action, but rather limiting conditions on how we act, also fails, because in requiring a

standing unconditional commitment to acting in accordance with impartial moral principles, it demands that the commitments underlying partial action be fragile. Both of these sets of attempts succeed in accommodating a kind of partial action, but both ultimately fail because they cannot reconcile impartial morality with the right kind of partiality – that is, the kind of partiality that is constitutive of authentic loving relationships.

The failure of the second set of attempts is particularly instructive, because it shows us that even if we can accommodate deep partiality within impartial moral theories, there is still a problem with those theories, insofar as they hold that all of the requirements generated by impartial moral principles are strict limiting conditions on our actions. This is because in doing so, they require an ordering of commitments in which the demands of impartial morality always override the demands of partiality to others (and, as I will discuss later, partiality to oneself).²¹ And while we do not expect that authentic loving relationships include unconditional commitments to our loved ones, we do, I think, expect that our commitments to them will not always take a back seat to the demands of impartial moral principles.

Overridingness and All-Things-Considered

If I am right that having an unconditional standing commitment to acting in accordance with impartial moral principles is incompatible with the appropriate level of commitment to partial action that is partially constitutive of authentic loving relationships, then we must give up one of three things: impartial morality, authentic loving relationships, or the thesis that we must be unconditionally committed to doing what morality demands. I take the latter thesis to be

²¹ I have borrowed the language of “ordering of commitments” from Susan Wolf (2012).

equivalent to the thesis that morality is always overriding. One way of stating the thesis that morality is overriding is to say that whenever a moral requirement comes into conflict with one or more non-moral requirements or prerogatives, the former overrides the latter. There is a question, of course, as to what exactly it amounts to for a moral requirement to override in this way. The answer to this question will depend on what one thinks about the relationship between various normative concepts like reasons, requirements, and “oughts”. But as I will conceive of it in this paper, for moral requirements to override amounts to the following: *If one is morally required to Φ , then, all things considered, one ought to Φ* . Now, if this is to be a substantive thesis, it must be the case that “all things considered, one ought to Φ ” amounts to something different from “one is morally required to Φ .” Thus, I will attempt to distinguish the meanings of these two statements and show that they do not amount to the same thing.

Let us start with the claim that one is morally required to Φ . To claim that one is morally required to Φ is to issue an overall moral verdict about Φ -ing. This claim should be distinguished, therefore, from the claim that one has some moral reason to Φ . To claim that one is morally required to Φ is to claim that, once all moral reasons have been taken into account, the set of such reasons taken as a whole rules decisively in favor of Φ -ing and against failing to Φ . Now, let us consider the claim that, all things considered, one ought to Φ . This also seems to be a verdictive claim; however, it isn’t a claim about the verdict that is issued about Φ -ing once all and only the *moral* reasons have been taken into account. Rather, it is a claim about that verdict that is issued about Φ -ing when *all relevant* reasons have been taken into account.²² Thus, whenever non-moral reasons are relevant, it is at least conceptually possible for what one ought

²² It may be the case that the all-things-considered verdict takes all relevant reasons into account directly, or it may be that it does so indirectly by taking into account all the verdicts (moral, etc.). Either way, it takes all relevant reasons into account, so I won’t take a position on this here.

to do all-things-considered to come apart from what one is morally required to do. The overridingness thesis, then, is the thesis that because a moral verdict can never be outweighed by non-moral reasons, what one ought to do all-things-considered will never come apart from what one is morally required to do.

My characterization raises the following question: from where is the verdict that takes all relevant reasons into account issued? When it comes to the verdict that takes all moral reasons into account, the answer seems clear: a moral verdict is being issued by morality; thus, we can point to morality as the source of moral verdicts. But it is less obvious what system issues the verdict that, taking all relevant reasons into account, one ought all-things-considered to Φ . What system could there be, over and above the system of morality, that takes both moral and non-moral considerations into account in this way? The real question here, I take it, is whether there is such a thing as an all-things-considered practical perspective. On this issue, I agree fully with Sarah Stroud:

Is there really a perspective which adjudicates between the reasons generated by all the various modes of evaluation which could be brought to bear on a practical predicament? Is it really possible to evaluate one's reasons as a whole, to weigh against each other the different kinds of reasons one has? Are reasons of different types generally *commensurable*? One reason for saying "yes" to these questions is this. When faced with a decision – when in a practical predicament – we must of course decide on *one* thing to do. Often (not always), in trying to make what we think is the *right* decision, we attempt to figure out what we have most reason to do in precisely this sense. We endeavor to survey and compare the full field of reasons, both for and against, with a view to determining which course of action is, all things considered, best supported by reasons. The aim of practical reasoning can thus very naturally be seen as that of arriving at a judgment as to what one has most reason to do (and then, of course, acting on it). Insofar as that seems to be what we are up to in our practical reasoning, we are committed to thinking that it makes sense to talk of what someone has most reason to do.²³

²³ Stroud, p. 175.

In short, it seems that the phenomenology of deliberation supports the claim that there is an all-things-considered practical perspective from which a verdict can be issued that takes into account all relevant reasons, both moral and non-moral.

Moreover, if there is not an all-things-considered practical perspective, I can see no way of making sense of the overridingness thesis. According to that thesis, if the moral verdict issued about Φ -ing comes into conflict with non-moral reasons, it always overrides them. From what practical perspective is this the case, if not from the all-things-considered one? It can't be simply that, from the moral perspective, if the moral verdict issued about Φ -ing comes into conflict with non-moral reasons, it always overrides them. The moral perspective takes into account all the moral reasons for and against Φ -ing and issues a moral verdict; thus, it seems it can't pronounce on whether moral verdicts are overriding, because to do so would be to take into account reasons that are outside its purview. So even if it is the case that moral requirements always override, if this claim is to be anything but trivial, it must be from an all-things-considered practical perspective that it is the case. But, as I have argued, if we are to reconcile the right kind of partiality with impartial moral theories, we must deny the overridingness thesis.

A Modest Defense of Impartiality

If I am right that the two attempts to reconcile partiality with impartial morality I discussed earlier both fail, then it seems to me that a successful internal defense of impartial morality from Stocker- and Williams-style objections does not seem to be in the offing. If, like me, one finds the idea of morality as an impartial system compelling, but also finds the objections compelling, one must turn to an external defense of impartial morality. By this, I

mean that instead of defending impartial morality at a normative ethical level, we must defend it at a meta-ethical level by saying something about how morality fits into the structure of practical reasoning. In my view, this meta-ethical defense must be a modest one, in the sense that it must make significant concessions about the practical authority of morality. According to the modest defense I will offer here, morality is indeed impartial, but is only one of (at least) three practical perspectives that come into play in our relationships with ourselves and others. Each of these practical perspectives, when relevant, issues verdicts about actions; my central claim is that none of these perspectives always issues verdicts that are overriding. Not only will this defense preserve the impartiality of morality, but it will also shed light on how the right kind of partiality is justified on non-moral grounds.

There is a long tradition in moral philosophy of thinking that impartiality is partly constitutive of morality – that part of what it is for a reason for action to be a distinctively *moral* reason is for it to be a reason that gets its normative force from an impartial general principle. For example, a utilitarian might think that what it is for someone to have a distinctively moral reason to help one's friends is for it to be the case that doing so would maximize happiness. The corresponding Kantian story would be a little bit more complicated, but it would presumably have something to do with the fact that the categorical imperative generates a perfect duty to help one's friends. I'm sympathetic to this picture of morality, because it seems to me to be an indispensable part of morality that every moral agent has equal standing. And what this equal standing amounts to, I think, is that for the purposes of what grounds our moral reasons for action, each person "counts" just as much as anyone else. Thus, each person is in an important

sense replaceable²⁴; if I have a moral reason to help her, I also have moral reasons to help any other similarly situated person. Thinking of morality as an impartial system in this way makes sense of our intuition that it is fundamentally a fair and equal system.

Thus, in an important sense, it seems to me that moral reasons, as the impartialist conceives of them, are impartial because they are impersonal, or, as I prefer to call them, *third-personal*. I call them third-personal because they do not get their normative force from special relationships we have with others, but rather just their membership in the community of moral agents. The lack of proximity or familiarity needed to address another second-personally (or first-personally) is what I think gives moral reasons a third-personal character. The moral perspective, then, is the practical perspective that takes into account all and only these third-personal reasons and issues moral verdicts on actions. I take this picture of the moral perspective to be compatible with both consequentialism and Kantianism, as well as other impartial moral theories like contractualism.

According to this picture, all moral reasons for action are grounded in impartial moral principles. Thus, any reasons we have that are grounded in facts about our special relationships with ourselves and others, unmediated by impartial general principles, are outside the purview of morality. In other words, the kinds of reasons which justify acting with *deep* partiality are not moral reasons. Thus, on my account, assuming some standard account of how reasons and obligations are connected, there are also no moral obligations that are fundamentally “special.”²⁵

²⁴ It may seem odd to claim that each person is replaceable according to Kantianism, given the centrality of the separateness of persons in Kant’s thought. However, I do think each person is replaceable in the sense that, for example, if the categorical imperative generates a duty for me to care for my child, it generates a duty for me to care for *anyone* who is my child.

²⁵ That is not to say, however, that there are no special moral obligations at all, if such obligations are to be conceived of as moral obligations to act partially toward certain other people. But just like they are obligations to act with surface partiality, these obligations are only special on the surface; ultimately, they are grounded in impartial moral principles, and thus fundamentally no more special than any other moral obligations.

Taking Stocker's and Williams' objections to impartial morality seriously means that we must admit that we have reasons to act partially toward loved ones such that, if we act on the basis of those reasons, we act with deep partiality. But according to this picture, these reasons cannot be moral reasons. Much of this paper has focused on how, for a defender of impartial morality, the kind of partiality that is constitutive of authentic loving relationships could be justified.

According to my account, the justification for such partiality comes not from morality at all, but from a different mode of evaluation altogether.

If both my wife and a stranger are drowning, all else being equal, most would say I am morally justified in saving my wife. Why is this? According to some of the impartialist views I have discussed, it is because, while I have moral reasons to save each, I have an additional moral reason to save my wife that is legitimated by some impartial general principle(s). It seems right to say I have this additional moral reason, but as Williams' "one thought too many" case shows, if this is the reason on the basis of which I save my wife, there seems to be something missing in our relationship. Williams writes of the man in his example that "it might have been hoped by some (for instance, by his wife) that his motivating thought, fully spelled out, would be the thought that it was his wife, not that it was his wife and that in situations of this kind it is permissible to save one's wife."²⁶ It seems that if I really deeply and wholeheartedly love my wife, it won't be for any moral reason that I act. Rather, it will be for a non-moral reason such as, "it's my wife," *unmediated* by any impartial general principle. If "it's my wife" were a moral reason, it would have to be grounded in the kind of principle gestured at by Williams, not the bare facts of my love for my wife. It is the facts of this special relationship that, unmediated by

²⁶ Williams (1981), p. 18.

any general principle, give me a reason to save her. This is not a moral reason, but a reason of love.

Thus, on my account, if I deeply and wholeheartedly love my wife, the reason on the basis of which I save her must be a reason of love, not a moral reason. Rather than being impersonal, this reason is deeply personal; my wife is in no sense replaceable, because my reason to save her gets its normative force from the facts of the unique, special relationship I have with her. This is, I think, an example of a distinctive kind of non-moral reason for action that I will call a second-personal reason. These are the reasons that justify the kind of partiality to others that is constitutive of authentic loving relationships. I call them second-personal reasons because in relying on our special relationships with others, they rely on a kind of proximity and familiarity that allows us to say to the other “I did it because it was *you*.” This is the kind of explanation my wife might hope for when she asks why I saved her, not some explanation that makes reference to impartial general principles. These are the reasons, I think, that those with whom we have a wide variety of special relationships hope we act for when we act partially toward them. If I am your close friend, and you need me to undertake some difficult task for you, the right kind of thing for me to say seems to be, “I wouldn’t do this for most people, but I’ll do it *because it’s you*.” Acting partially toward others for these “it’s you” kind of reasons, which I am calling second-personal reasons, seems to me to be constitutive of relating to friends, family and partners in the right way. What we want from our loved ones is for them to relate to us from this second-personal perspective²⁷, acting partially toward us based on its verdicts.

²⁷ What I am calling the “second-personal perspective” is not in any relevant sense related to Darwall’s notion of the “second-person standpoint” (see Darwall (2006)). This is mostly an unfortunate coincidence of terminology; but it is true that my analysis of moral reasons as third-personal in character is at odds with his analysis of them as second-personal in character.

I have now discussed two of the three practical perspectives that figure into my account: the third-personal, and the second-personal. The second-personal perspective is the most important in responding to Stocker's and Williams' worries, because it provides an account of the justification of the kind of partiality to others that is constitutive of authentic loving relationships. However, there are also corresponding worries, which I have not discussed so far, that impartial morality is also incompatible with the right kind of partiality to *oneself* that might be thought to be constitutive of a healthy self-conception. Williams himself raises this worry when he discusses consequentialism's incompatibility with the pursuit of ground projects. He argues that part of the problem with the impartiality of consequentialism is that " 'it's me' can never in itself be a morally comprehensible reason."²⁸ I agree with Williams that just as we need to be able to make room for reasons of the form "it's you," in our practical lives, we need to be able to make room for reasons of the form "it's me." However, I think the impartialist can grant this while maintaining that these reasons are not morally comprehensible (or at least not within the purview of morality).

When I decide to be partial to myself for the reason that it's me, this is also not a moral reason, but a reason of another kind. I will call these reasons of partiality to oneself *first-personal* reasons, because of the role "it's me" plays in them. Reasons of partiality to oneself are often described as prudential reasons or reasons of self-interest; however, I think it is a mistake to identify first-personal reasons with prudential reasons, as the former category is somewhat broader than the latter. What I do for myself is not necessarily restricted to what will make me happiest or most well-off. For example, it seems plausible to say that, for some ground projects, we have first-personal reasons to engage in them. But it seems overly narrow to describe it as out

²⁸ Williams (1973), p. 32.

of prudence of self-interest that we do so. Nevertheless, prudential reasons will make up a large subset of first-personal reasons. Much of what I will say about first-personal reasons will parallel what I have said about second-personal reasons, treating partiality to the self analogously to partiality to others. In this regard, acting on the basis of first-personal reasons can be said to be constitutive of a kind of self-love.

Consider, for example, a modified version of Williams' case, in which I must choose not between saving my wife and saving a stranger, but between saving myself and saving a stranger. Similarly, I think I have moral reasons both to save myself and the stranger, and I may even have an additional moral reason to save myself. But I think I also have an additional first-personal reason to save myself. If I save myself on the basis of my additional moral reason, it is not out of self-love that I save myself. But if I save myself primarily on the basis of my first-personal reason, I act with deep partiality toward myself. Thus, it does seem to be out of a kind of self-love that I do so. Whether we call it self-love, a healthy self-conception, or something else, it seems to me that like truly and deeply caring for our friends and loved ones requires inhabiting the second-personal perspective, truly and deeply and deeply caring for ourselves requires inhabiting the first-personal perspective.

The Ordering of Practical Perspectives

I have suggested that in order to reconcile impartial morality with the kind of partiality to others that is constitutive of authentic loving relationships, we must concede that morality is not overriding. On the account I have offered, morality is but one of at least three practical perspectives that come into play in our relationships with ourselves and others. I have called

these practical perspectives the first-, second-, and third-personal perspectives. By a practical perspective, I simply mean the evaluative stance from which a verdict is issued about an action, taking into account all the relevant reasons of the appropriate kind. I have also argued that there is an all-things-considered perspective from which a verdict is issued about an action, taking into account all the relevant reasons *simpliciter*. It is this verdict that tells us what, all things considered, we ought to do. If, as I've argued, morality is not always overriding, then it is possible, at least in principle, that sometimes what we ought to do all-things-considered will come apart from what we are morally required to do. In other words, we can do what we ought to do all-things-considered without having the standing unconditional commitment to morality, and thus to impartiality, that gets in the way of having authentic loving relationships. Importantly, to say that morality is not always overriding is not to say that morality is always *overridden*. To claim this would be to carve out for morality an implausibly weak role in our practical lives. Instead, on my account, no practical perspective always overrides, nor is any always overridden.

Insofar as these practical perspectives come into conflict, our practical lives are a balancing act between them. Positing an all-things-considered practical perspective makes room for an account of how we performing this balancing act. Of course, simply positing this all-things-considered perspective does not provide us with such an account; how exactly the all-things-considered perspective weighs reasons of various kinds and generates authoritative verdicts about actions is still an open question. But however that account goes, positing an all-things-considered perspective from which morality is not always overriding does make room for love, because if that is the case, there is no requirement of practical reason that an agent must have a standing unconditional commitment to morality. Unfortunately, it is beyond the scope of

this paper to defend my own account of how all-things-considered verdicts are issued; however, I will say one important thing about how I think it such an account would go.

My goal has been to show that we must make room for love in the sense that it must be open to agents to have the kind of commitments constitutive of authentic loving relationships, and I have argued that those commitments must be robust enough that they are incompatible with a standing unconditional commitment to impartial morality. However, it would be odd to claim that if an agent did have a standing unconditional commitment to impartial morality, she would necessarily be irrational, or in some other way deficient from the all-things-considered perspective. She would be, as I have argued, unable to engage in authentic loving relationships. But if some people are interested in being, to borrow from Susan Wolf, “moral saints,”²⁹ even at the expense of such relationships, I see no reason to claim that they are not doing what they ought to do all-things-considered. But most of us *are* interested in engaging in authentic loving relationships, and so as long as we have *some* appropriate level of commitment to morality, I also see no reason to think we are not doing what we ought to do all-things-considered, even if it means we are not unconditionally committed to morality.

Thus, while I cannot offer a complete account of it here, I can say that I am attracted to a kind of *permissivism* about the all-things-considered practical perspective. It might be that for agents who really are unconditionally committed to morality, moral requirements will, in practice, always override. But for the rest of us, morality will only sometimes be overriding. I think it is quite intuitively plausible to think that the all-things-considered perspective is permissive such that there are multiple permissive orderings of commitments. If this is the case, there must nevertheless be some constraints on which orderings are permissible; for example, it

²⁹ Wolf (1982).

seems obviously impermissible to have no commitment to morality at all. Again, it is beyond the scope of this paper to give a full account of these putative constraints. But as long as such an account makes room for the kind of ordering of commitment I have argued is necessary for authentic loving relationships, it will make room for love.

What I have offered above may seem to render practical reason hopelessly messy, but I think it is actually quite neat, considering the richness and complexity of our practical lives. Moreover, unless my own practical deliberations are quite eccentric, it seems to me that the categories of practical reason I have posited are more true to the phenomenology of practical deliberation than other accounts have been. On my account, impartial morality is not inconsistent with the kind of partiality that is constitutive of authentic loving relationships. Impartial moral requirements and the requirements of deep, robust partiality, simply arise out of different practical perspectives that factor into the all-things-considered verdict on what we ought to do. Most of us are interested both in being morally good and in having authentic loving relationships. This, I have argued, requires that we be unconditionally committed neither to morality nor to our loved ones. As Susan Wolf succinctly puts it, in bringing home her interpretation of Williams, either of these commitments would be “one unconditional commitment too many.”³⁰ Thus, we must be open to the possibility that love and morality will come into deep conflict, and neither will always win out. But for many of us, I think, that is just life.

³⁰ Wolf (2012), p. 92.

Conclusion

In this paper, I have sought to show that the conflict between impartial morality and the kind of partiality necessary for love is deep and complex, and to provide a modest defense of impartial morality on which the two can coexist. As I have argued, this kind of partiality must not only be deep, in the sense of being fundamentally partial in its justification, but also robust, in the sense of not being unconditionally overridden by the demands of impartial moral principles. In order for this kind of partiality to coexist with impartial morality, it must be the case that morality is not always overriding. On the account I have offered, morality is impartial, but is one of multiple practical perspectives, none of which is necessarily overriding, that factor into the all-things-considered practical perspective. I find this picture of our practical lives not only to be independently plausible, but also to be the best response to the critiques of impartial morality raised by Stocker and Williams. It may be that in attempting to reconcile impartial morality with the kind of partiality that is constitutive of authentic loving relationships, I have provided a theory that is unsatisfying to those on both sides of the debate. It certainly will not be satisfying to those with the strong conviction that morality *is* overriding, always and for everyone. But for those like myself, who do not have this conviction, but do have the conviction that morality is impartial, I hope the beginnings of an account I have offered here will seem to be an attractive and plausible solution.

REFERENCES

- Baron, M. (1984). The Alleged Moral Repugnance of Acting from Duty. *The Journal of Philosophy*, 81(4), 197-220.
- Baron, M. (1991). Impartiality and Friendship. *Ethics*, 101(4), 836-857.
- Darwall, S. L. (2006). *The second-person standpoint: Morality, respect, and accountability*. Cambridge, MA: Harvard University Press.
- Hill, T. E. (1991). *Autonomy and self-respect*. Cambridge: Cambridge University Press.
- Kant, I. (1996). *The metaphysics of morals* (M. J. Gregor, Trans.). Cambridge: Cambridge University Press.
- Railton, P. (1984). Alienation, Consequentialism, and the Demands of Morality (1984). *Philosophy and Public Affairs*, 134-171.
- Sidgwick, H. (1907). *The methods of ethics* (7th ed.). New York: MacMillan.
- Smart, J. J., & Williams, B. (1973). *Utilitarianism: For and against*. Cambridge: Cambridge University Press.
- Stocker, M. (1976). The Schizophrenia of Modern Ethical Theories. *The Journal of Philosophy*, 73(14), 453-466.
- Stroud, S. (1998). Moral Overridingness and Moral Theory. *Pacific Philosophical Quarterly*, 79(2), 170-189.
- Williams, B. (1981). *Moral luck: Philosophical papers, 1973-1980*. Cambridge: Cambridge University Press.
- Wolf, S. (1982). Moral Saints. *The Journal of Philosophy*, 79(8), 419-439.
- Wolf, S. (2012). 'One Thought Too Many': Love, Morality, and the Ordering of Commitment. Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams, 71-92.