UC San Diego
UC San Diego Electronic Theses and Dissertations

Title

The metaphysics of morality : a dispositionalist account

Permalink

https://escholarship.org/uc/item/6c24r6tc

Author

Robinson, William Luke

Publication Date

2005

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Metaphysics of Morality: A Dispositionalist Account

A dissertation submitted in partial satisfaction of the

requirements for the degree of Doctor of Philosophy

in

Philosophy

by

William Luke Robinson

Committee in charge:

   Professor David O. Brink, Chair
   Professor Richard J. Arneson
   Professor Anthony Edwards
   Professor Alan C. Houston
   Professor Samuel C. Rickless

2005

The dissertation of William Luke Robinson is approved,

and it is acceptable in quality and form for publication on

microfilm:

Richard J Arneson

Samuel C Rickless

Anthony 2 Edwards

Ala

David O. Brink

Chair

University of California, San Diego

2005

iii

To my dad

It is evident that we are at a primitive stage of moral development. Even the most civilized human beings have only a haphazard understanding of how to live, how to treat others, how to organize their societies. The idea that the basic principles of morality are *known*, and that the problems all come in their interpretation and application, is one of the most fantastic conceits to which our conceited species has been drawn. (The idea that if we cannot easily know it, there is no truth here is no less conceited.) Not all of our ignorance in these areas is ethical, but a lot of it is. And the idea of the possibility of moral progress is an essential condition of moral progress. None of it is inevitable.

Thomas Nagel, *The View From Nowhere*

If nothing further were at stake, I should not be particularly concerned about whether we believe in a ruly world or in an unruly one, for, not prizing the purity of our affirmations, I am not afraid that we might hold false beliefs. The problem is that our beliefs about the structure of the world go hand-in-hand with the methodologies we adopt to study it. The worry is not so much that we will adopt wrong images with which to represent the world, but rather that we will choose wrong tools with which to change it. We yearn for a better, cleaner, more orderly world than one that, to all appearances, we inhabit. But it will not do to base our methods on our wishes. We had better choose the most probable option and whenever possible hedge our bets.

Nancy Cartwright, *The Dappled World*

TABLE OF CONTENTS

VITA

1992	B.A. *summa cum laude* in History, Auburn University
1996	J.D. *cum laude*, University of Pennsylvania Law School
1996-1998	Attorney, Philadelphia, PA
1998-2005	Teaching Assistant, University of California, San Diego
2001	M.A. in Philosophy, University of California, San Diego
2005	Ph.D. in Philosophy, University of California, San Diego

PUBLICATIONS

"Moral Holism, Moral Generalism, and Moral Dispositionalism," *Mind* (forthcoming)

ACKNOWLEDGEMENTS

ABSTRACT OF THE DISSERTATION

The Metaphysics of Morality: A Dispositionalist Account

by

William Luke Robinson

Doctor of Philosophy in Philosophy

University of California, San Diego, 2005

Professor David O. Brink, Chair

Each of the following theses is independently plausible:

Principle: Morality is principled: right-making factors or moral reasons entail corresponding moral principles.

Conflict: Conflicts of moral obligation are possible and metaphysically unproblematic (even if they can be quite difficult—or even impossible—to resolve).

Holism: Whether a given factor contributes to the rightness or wrongness of an action may depend on factors other than itself; hence, what is a right-making factor or moral reason to ϕ in one case may not be one in another, or may even be a wrong-making factor or moral reason not to ϕ given suitable circumstances.

Nevertheless, neither Conflict nor Holism has fared particularly well in ethical theory, perhaps because neither is obviously consistent with Principle. Moreover, on one influential account of moral principles, Conflict and Holism would seem to be inconsistent with Principle. On this account, moral principles state what Ross calls "*prima facie* obligations". But if, as Ross appears to hold, *prima facie* obligations are not genuine obligations and genuine obligations cannot conflict, then Conflict is inconsistent with Principle. And if, as many hold, *prima facie* obligations correspond to factors that are

always right-making factors or moral reasons to ϕ, then Holism is inconsistent with Principle.

An adequate moral metaphysic would reconcile these three theses: it would both show them to be mutually consistent and explain them. An adequate moral metaphysic would also explain why moral principles are capable of playing the prominent and important roles they play in moral thought, practice, and theory.

In this dissertation, I develop a moral metaphysic that takes moral dispositions or powers to be the fundamental units of moral explanation—that is, a moral-theoretic *dispositionalism.* On this view, right-making factors just are factors that have the power to make actions (*prima facie*) obligatory, and moral principles ascribe moral dispositions to *sorts* or *kinds* of actions, states of affairs, beings, institutions, and the like. This view explains why moral principles can play the roles they do in moral thought, practice, and theory. It shows Principle, Conflict, and Holism to be mutually consistent. And it explains each of them.

INTRODUCTION

If some factor makes acting in a certain way right, or is a moral reason to act in that way, then there must be a moral principle to that effect. For example, if promising to do something is a moral reason to do it, then 'One ought to keep one's promises', or something very much like it, must be a (true) moral principle. Such moral principles can conflict in particular cases, and one can have conflicting obligations. One might, for example, be obligated both to keep a promise to meet a friend and to stop and assist an injured person, even though one cannot do both. Normally, promising to do something obligates one to do it, and normally the fact that one made a promise continues to obligate one to keep it, even when other factors (such as someone's being injured) obligate one to do other things that are incompatible with doing so (such as stopping to help an injured person). But suppose that the promise in question is to meet a friend for some nefarious purpose. Would the fact that one promised to do something—to meet one's friend—obligate one to do it in such a case? Would it be *any* moral reason to do it, or make doing it even *pro tanto* right? Might it not even obligate one *not* to do it, or be a moral reason *not* to do so? It seems that whether the fact that one promised to ϕ is a moral reason to ϕ, no moral reason to ϕ, or a moral reason not to ϕ, depends—or at least may depend—on factors other than itself, such as the nature or purpose of the promised action.

Claims such as these express three distinct theses about morality:

Principle: Morality is principled: right-making factors or moral reasons entail corresponding moral principles.

Conflict: Conflicts of moral obligation are possible and metaphysically unproblematic (even if they can be quite difficult—or even impossible—to resolve).

Holism: Whether a given factor contributes to the rightness or wrongness of an action may depend on factors other than itself; hence, what is a right-making factor or moral reason to φ in one case may not be one in another, or may even be a wrong-making factor or moral reason not to φ given suitable circumstances.

Each of these theses is independently plausible, and each has its advocates, but they have not fared equally well in ethical theory. Moreover, some have found them mutually inconsistent.

Principle appears to be widely held, even more widely assumed, and rarely questioned. Kant defines "morals" as the science of the laws in accordance with which everything ought to happen (Kant 1785, 4: 387-88). Mill takes it as uncontroversial that "morality must be deduced from principles," and "that the morality of an individual action is not a question of direct perception, but of the application of a law to an individual case" (Mill 1871, 1.3). And, more recently, Scanlon claims that "to justify an action to others is to offer reasons supporting it and to claim that they are sufficient to defeat any objections that others may have," and that "to do this . . . is also to defend a principle, namely one claiming that such reasons are sufficient grounds for so acting under the prevailing conditions" (Scanlon 1998, 197).

Neither Conflict nor Holism enjoys anything like the wide acceptance that Principle does. Kant thought conflicts of obligation were logically impossible. For him, "*obligation*

is the necessity of a free action under a categorical imperative of reason" and "[a] *duty* is that action to which someone is bound" (Kant 1797a, 6: 222). As two incompatible actions cannot both be necessary under a categorical imperative, conflicts between duties or obligations are, *for Kant*, inconceivable. The most Kant can conceive are conflicts between "grounds of obligation" (6: 224). But this will strike those who find Conflict plausible as absurd. Kant is certainly free to define *his* concepts, but he cannot derive a substantive conclusion from those concepts without begging the question of whether they are, or should be, *our* concepts.

Mill could at least conceive of conflicts of obligation. "There exists," he writes, "no moral system under which there do not arise unequivocal cases of conflicting obligation" (Mill 1871, 2.25). However, it is not clear just *how* Mill conceives of such conflicts. Mill is often read as an act-utilitarian. Act-utilitarians are often said to hold that there is just one obligation: to produce the greatest possible balance of happiness over unhappiness. And that (supposed) obligation seems incapable of conflicting with itself. Thus, one might wonder whether there are any unequivocal cases of conflicting obligation in Mill's own moral system. It could be that Mill thinks conflicts of obligation are conceivable, yet impossible in some weaker sense; but this is inconsistent with what he says. Another possibility is that he conceives of conflicts of obligation subjectively, as situations in which it is difficult to discern what one's obligation is, rather than objectively, as situations in which one has conflicting obligations. Yet a third possibility is that he conceives of conflicts of obligation objectively and thinks that they do arise. Unfortunately, whatever his view, Mill offers us no clear answer to the question of what

such unequivocal cases of conflicting obligation consist in. That is, he offers us no explicit account of the metaphysics of conflicts of obligation.

Unlike Mill, Ross does offer an explicit account of the metaphysics of conflicts of obligation (Ross 1939, 83-86; Ross 1930, 17-20). Unfortunately, it appears to be an eliminativist one, for Ross appears to hold that conflicts of obligation are impossible. On Ross's account, conflicts of obligation are conflicts between what he calls "*prima facie* obligations", rather than conflicts between what he calls "obligations", and conflicts between the latter are impossible. Moreover, it is standard to gloss *prima facie* obligations as, or in terms of, moral reasons or right-making factors. And conflicts between moral reasons and right-making factors are not obviously conflicts between obligations. Indeed, they would seem to be conflicts between the grounds of obligation.

If Conflict has fared poorly in ethical theory, Holism, or *moral holism*, has fared even worse. Indeed, it seems that no one even thought to raise the question of its merits until quite recently. Moral holism claims that the moral polarity or "valence", as it were, of a factor may be affected by other factors, and so need not be the same in every case. Moore implicitly denies that such a thing is possible when he claims that "'right' does and can mean nothing but 'cause of a good result', and is thus identical with 'useful'" (Moore 1903, 147). If that were true, then not only would it make no sense to say that something other than an action's usefulness might contribute to its rightness, it would make no sense to say that an action's usefulness might fail to contribute to its rightness. Thus, Moore's view implies that the holist's thesis is, literally, nonsensical. Hare, too, implicitly denies that such a thing is possible when he claims that if one says that an action is right and thinks that one is entitled to do so because it is F, then one is committed to the view

that every F action is right (Hare 1963, 10-12). If that were true, one could not

consistently claim that one action is right simply because one promised to perform it but

that a second is not even *pro tanto* right because, although one promised to perform it,

that fact makes no contribution to the rightness of doing so because the action would be

a nefarious one. Thus, Hare's view implies that moral holism is logically false.

If moral holism is sometimes implicitly denied, its contradictory, *moral atomism*, is

(perhaps) more often implicitly assumed. It is, for instance, presupposed by many

familiar, seemingly cogent moral arguments. One is James Rachels's well-known

argument against the view that killing someone is, in itself, morally worse than letting her

die (1975, 228-29). Rachels asks us to consider "two cases . . . that are exactly alike except

that one involves killing whereas the other involves letting someone die" (228), the case

of Smith and the case of Jones:

> Smith stands to gain a large inheritance if anything should happen to his
> six-year-old cousin. One evening while the child is taking his bath. Smith
> sneaks into the bathroom and drowns the child, and then arranges things
> so that it will look like an accident. . . . Jones also stands to gain if
> anything should happen to his six-year-old cousin. Like Smith, Jones
> sneaks in planning to drown the child in his bath. However, just as he
> enters the bathroom Jones sees the child slip and hit his head, and fall face
> down in the water. Jones is delighted: he stands by, ready to push the
> child's head back under if it is necessary, but it is not. With only a little
> thrashing about, the child drowns all by himself, "accidentally," as Jones
> watches and does nothing. (228-29)

Rachels judges that there is no morally significant difference between Smith's action and

Jones's omission—that they are on a moral par. That much is consistent with moral

holism. But Rachels also claims that if these cases are on a moral par, then "the bare

difference between killing and letting die does not, in itself, make a moral difference"

(229)—by which he means not just that this bare difference does not, in itself, make a

moral difference in the case of Smith (or between it and the case of Jones), but that this bare difference does not, in itself, make a moral difference in any morally possible case (or between any two such cases).[1] From (*a*) 'That ϕ-ing would be a killing rather than a letting-die is morally indifferent in the case of Smith', Rachels simply *deduces* (*b*) 'That ϕ-ing would be a killing rather than a letting-die is, in itself, morally indifferent in any case—i.e., that it is an invariantly morally indifferent factor'. In doing so, he implicitly assumes that the moral valence of a factor *f* could not vary between two cases, and that the moral valence of a factor *f* could not be affected by factors other than *f* itself. In other words, his argument presupposes moral atomism.

Why have Conflict and Holism fared so poorly in ethical theory? Perhaps it is because neither is obviously consistent with Principle. If the morality of an individual action is a question of the application of a law to an individual case (as Mill says), and if moral principles require or forbid performing certain kinds of action categorically (as many, including Kant and Hare, have thought), then conflicts of obligation are logically impossible. Many now deny that moral principles require or forbid performing certain kinds of actions categorically and hold with Ross that they state *prima facie* obligations. But if *prima facie* obligations are not obligations (as Ross appears to hold), then this account of moral principles does nothing to reconcile Conflict with Principle. Moreover, received accounts of what *prima facie* obligations are identify or associate them with factors that are *always* right-making factors or moral reasons to ϕ. Thus, we get the view

---

[1] This is clear from his claim that the lack of a moral difference between the case of Smith and the case of Jones shows that active euthanasia is not, in itself, worse than passive euthanasia *even though the two sorts of cases are nothing alike* (Rachels 1975, 229).

that moral principles are or entail *invariant moral reasons*—factors whose moral valences are invariant across all contexts. If this is right, then Holism is inconsistent with Principle, for principle says that moral reasons entail corresponding moral principles. If moral reasons entail corresponding moral principles, and if moral principles entail invariant moral reasons, then—assuming that the moral reasons that entail moral principles are the same invariant moral reasons that those principles entail—a factor's moral valence must be invariant across all contexts. But that is just what Holism denies.

If Conflict and Holism are inconsistent with Principle, it does not follow that they are false. They might be true and Principle false. Thus, one might deny Principle on the basis of one or both of Conflict and Holism. I am not aware of anyone who denies Principle on the basis of Conflict, but some "moral particularists"—most notably Jonathan Dancy—evidently reject Principle on the basis of Holism. They maintain that moral holism motivates skepticism about the existence of and need for moral principles, along with skepticism about the viability of principle-based approaches to ethics and moral theory. But if each of these three theses—Principle, Conflict, and Holism—is independently plausible, then it is reasonable to seek a moral metaphysic—a metaphysics of, among other things, moral principles and right-making factors—on which they are all true. Ideally, such a metaphysic would not just be consistent with these three theses but also explain them.

One goal of this dissertation is to develop such a moral metaphysic. But its principal goal is to justify the view that moral principles are indeed capable of playing the prominent and important roles they play in diverse areas of moral thought, practice, and theory. Much moral and other practical thought is "intuitive", or at least not consciously

deliberative or inferential. Nevertheless, we appeal to moral principles in various areas of moral thought and practice. We appeal to them in, among other things, moral deliberation and advising, moral justification and criticism, and moral persuasion and argument. They play prominent and important roles in moral learning and education. We regard them as having normative force, and, in some cases, we try to live by them. They also play a central role in moral explanation—in explaining moral phenomena such as the rightness of particular actions, the justness of particular social systems and institutions, and the nature and extent of moral rights.

Unlike some who maintain the existence of moral principles but seem unconcerned with whether or not they can or do play any of these roles (e.g., Kagan 1998, 184-89; Jackson, Pettit, and Smith 2000), I aim to defend the claim that moral principles can indeed play them. I also aim to explain why this is so, and why moral particularists are wrong to think that moral holism motivates skepticism about moral principles and principle-based approaches to ethics and moral theory.[2]

---

[2] Even Dancy, who more than any other particularist wants to have nothing to do with moral principles, admits: "Particularists . . . face an uphill struggle, because principled conceptions of the ethical are culturally dominant and have been for centuries, at least in the West. And it is not as if nobody acts in accordance with such conceptions. People do 'have' principles, and they do (sometimes, at least) appeal to them in making their own decisions and in judging the actions of others. They condemn anyone who fails to act on his own principles, or those whose principles they take not to be the right ones. And they try to act 'consistently', where this means making their decisions in such a way that they can all be fitted under the same set of principles. So it is not as if all this talk of principles is mere theory. People actually order their lives (and the lives of others) according to their principles, and it is not very surprising that they are reluctant to admit that there is any chance the whole thing is a mistake" (Dancy 2004, 2). He also admits: "It is the job of a philosopher, so far as possible, to give an account of our practice rather than to tell us that we all ought to be doing something else. To the extent this cannot be done, it is normally a fault in the philosophy rather than in the practice. So particularism needs to offer some account, within the constraints which it accepts, of what is a very common practice of somehow appealing to general truths and previous cases in the course of reaching a moral judgment, and in the justification of one when reached" (Dancy 1993, 67).

I begin by sketching an account of the metaphysics of morality, which I dub *moral generalism*, and by arguing that this account explains why moral principles play the roles they do and vindicates the belief that they can do so. On this view, *general moral facts*—facts about the moral attributes of kinds (the properties of kinds of actions and states of affairs, the moral rights of kinds of beings, etc.)—explain *particular moral facts*—facts about the moral attributes of particulars (the moral properties of particular actions and states of affairs, the moral rights of particular beings, etc.), and true moral principles state or express general moral facts. According to moral generalism, moral principles can and do play the roles they play in moral thought, practice, and theory because true moral principles are statements of general moral fact and general moral facts explain particular moral facts. Ultimately, I advance a moral metaphysic that takes dispositions or powers to be the fundamental units of moral explanation—that is, a moral theoretic version of *dispositionalism*. On this view, right-making factors just are factors that have the power (or, to use Ross's term, a tendency) to make actions *prima facie* obligatory, and moral principles ascribe moral dispositions (powers, tendencies, etc.) to *sorts* or *kinds* of actions, states of affairs, beings, institutions, and the like. Insofar as *moral dispositionalism* explains moral generalism, it also explains why moral principles play the prominent and important roles they do in moral thought, practice, and theory—and why they are capable of doing so. Furthermore, it explains Principle, Conflict, and Holism.

CHAPTER 1

MORAL GENERALISM

Moral generalism is the view that *general moral facts* explain *particular moral facts*, and that true moral principles state or express general moral facts. (Alternately, moral principles are general moral facts, and true *statements* of moral principle state or express such facts.) By general moral facts, I mean facts about the moral attributes of *kinds*—facts about the moral properties of kinds of actions (e.g., that homicide is wrong) and states of affairs (e.g., that unequal distributions of political rights are unjust), the moral rights of kinds of beings (e.g., that humans have the right to life, liberty and security of person), the moral properties of kinds of governments and economies (e.g., that liberal constitutional democracies are just and that slave economies are unjust), and so on.[1] And by particular moral facts, I mean facts about the moral attributes of particulars—facts about the moral properties of particular actions and states of affairs, the moral rights of particular beings, the moral properties of particular governments and economies, and so on.

Moral generalism, including its account of moral principles, is suggested by reflection on (some of) the kinds of moral judgments we make and the relations between and among

_____

[1] I do not mean, as is sometimes meant by the term 'general facts', *universal* facts or truths—those expressed by true universally quantified statements about particulars and their properties (e.g., that *all* homicides are wrong). Such facts are about classes of individuals, rather than kinds, and are at the same level of explanation as particular moral facts. Nor do I mean 'non-specific facts'.

these judgments and their objects. It also accords well with the prominent and important roles that appeals to *general moral considerations*—moral principles and (putative) general moral facts—play in diverse areas of moral thought, practice, and theory. Indeed, it explains this phenomenon. Moral generalism explains why we appeal to general moral considerations in, among other things, moral deliberation and advising, moral justification and criticism, and moral discourse and argument. It explains why general moral considerations play prominent roles is moral learning and education. It explains why we regard them as having normative force and—sometimes—strive to live by them. And it explains why moral principles play a central role in moral explanation—in explaining moral phenomena such as the rightness of particular actions, the justness of particular social systems and institutions, and the nature and extent of moral rights.

1. Moral Judgments and Moral Facts

Among the kinds of moral judgments we make are *particular moral judgments* and *general moral judgments*. Suppose that while hiking in the mountains, I come across Abby, who is clearly injured. The judgment that I ought to stop and help her is a particular moral judgment. It is a judgment about what a particular agent ought to do. The judgment that my not stopping to help Abby would be wrong is also such a judgment. It is a judgment about a particular action or omission—namely, my not stopping to help her. In contrast, the judgment that we ought to help others in distress is a general moral judgment, as is the judgment that hikers ought to help other hikers if they are injured. Each of these judgments is one about what a kind of being ought to do. The first is a judgment about what agents ought to do, while the second is a judgment about what

hikers ought to do.[2] Each is also a judgment about what *kind* of thing a certain kind of

being ought to do: they are judgments about the general obligations of being kinds.

We make particular moral judgments about numerous things, including the morality

or moral properties of particular actions and practices, the moral properties of particular

institutions, and the moral rights and obligations of particular people. The judgment that

the Third Reich's practice of systematically killing those it deemed undesirable was grossly

immoral is another example of such a judgment. So too are the judgment that the current

German government is a just one, the judgment that the slave economy of the

Antebellum South was an unjust one, the judgment that Anne Frank had a right to life

and Dred Scott a right to liberty, and the judgment that I have an obligation to stop and

help Abby. Such judgments are judgments about the moral attributes of particulars (their

morality, moral properties, rights, obligations, etc.).

We make general moral judgments about numerous things, too, including the

morality of acting in kinds of ways, the moral properties of kinds of actions, practices,

and institutions, and the moral rights and obligations of kinds of beings. The judgment

that we ought to help others in distress and the judgment that hikers ought to help other

hikers if they are injured are but two examples. The judgment that one ought to respect

the religious, moral, and cultural beliefs and traditions of others is another one. So too

are the judgments that slavery and genocide are wrong, that democracies are just and

dictatorships unjust, that humans have "the right to life, liberty and security of person"

---

[2] There is room to dispute what the relevant kinds are here, just as there is room to dispute the truth of such judgments. For example, one might argue that the former judgment is one about what people or humans ought to do, rather than one about what agents ought to do. Such matters are, however, orthogonal to the general character of such judgments.

(Universal Declaration of Human Rights, Art. 3), that women have the right to abort an

unwanted or dangerous pregnancy, and that a citizen has an obligation to fight for his or

her country in time of war.[3] Such judgments are judgments about the moral attributes of

kinds, or *genera.*

Moral judgments can, of course, be quite controversial. Few would deny that

murder and genocide are wrong, but people can and do disagree about many moral

matters. That is, people can and do disagree about what the facts of the matter are on a

range of moral questions. This is only to be expected. For moral judgments are

judgments about facts, moral facts, and just as it is it is not always clear what the non-

moral facts are, it is not always clear what the moral facts are.[4] Nevertheless, we can say

some illuminating things about moral judgments and moral facts as such, quite

independently of what the moral facts of the matter on specific moral questions are.

_____

[3] I speak hear of moral rights and obligations, rather than legal ones. The Universal Declaration of Human Rights recognizes human rights and fundamental freedoms, and aims to promote universal respect for and observance of them. It does not purport to create or confer them. Article 3 declares, "Everyone has the right to life, liberty and security of person." (This is not a general moral judgment in my sense—on which, see below.) Inasmuch as the Declaration recognizes rights and does not purport to create or confirm them, and inasmuch as the legal protection of the rights it recognizes is one of the things the Declaration aims to promote, those rights must be moral rights. Similarly, many of those who advocate womens' rights advocate the legal recognition and protection of what they believe to be moral rights. Indeed, one might well believe that governments have a moral obligation to protect the moral rights of their citizens, and of others within their jurisdiction. And a patriot may certainly believe that she has an obligation to defend her country even if she has no legal obligation to do so.

[4] These are hardly uncontroversial claims. The existence of moral facts is controversial, as is the prior question of whether moral judgments are in fact judgments (i.e., about facts). I am not here concerned to defend either moral cognitivism or moral realism. I take it that existing defenses of them (e.g., Brink 1989) are sufficiently robust to justify exploring other questions of interest to moral realists. That said, should such an exploration reveal features of moral thought and practice that non-cognitivists or other anti-realists have difficulty accounting for, so much the worse for anti-realism.

As judgments about the moral attributes of particulars, particular moral judgments are correct just insofar as they correspond to *particular moral facts*, facts about the moral attributes of particular actions, practices, institutions, beings, and so on. These include facts about the morality or moral properties of particular actions and practices. That my stopping and helping Abby would be right would be such a fact, as would the fact that the Third Reich's practice of systematically killing those it deemed undesirable was grossly immoral. They also include facts about the moral properties of particular institutions, such as the fact (if it is a fact) that the current German government is just, and the fact that the slave economy of the Antebellum South was unjust. And they include facts about the moral rights and obligations of particular beings. That Anne Frank had a right to life and that I have an obligation to stop and help Abby would both be particular moral facts.

Similarly, general moral judgments are correct just insofar as they correspond to *general moral facts*, facts about the moral attributes of kinds of actions, practices, institutions, beings, and so on. The fact that helpings of others in distress are right would be such a fact. So, too, would be the fact that genocide is wrong, the fact that democracies are just, the fact that humans have a right to life, and the fact that a citizen has an obligation to fight for his or her country in time of war. Such facts must be distinguished from *universal moral facts*, which are facts about classes and the members of which they consist, such as that all the members of a certain class (e.g., the class of human beings) possess a certain moral attribute (e.g., the right to life), or that two classes (e.g., right acts and maximally utile acts)—at least one of which is a moral class—are

coextensive. General moral facts are facts about kinds, which, although their members

do form classes, are not themselves classes.[5]

Evidently, general moral facts explain particular moral facts. If I ought to stop and

help Abby (a particular moral fact), for example, it is because of some general moral

fact—plausibly, that one ought to help others in distress. In other words, that one ought

to help others in distress, or some such general moral fact, explains why I ought to stop

and help Abby. This is not to deny that I ought to help Abby because she is in distress.

In fact, her being in distress seems not only the most obvious reason to stop and help

her, but also the most obvious reason for why I ought to stop and help her.[6] But if the

fact that Abby is in distress explains why I ought to stop and help her, it must be because

of something general—that one ought to help others in distress, that helping others in

distress tends to promote utility, that another's being in distress is a right-making factor,[7]

---

[5] We do sometimes judge that universal moral facts obtain. But we are far more reluctant to make such judgments than we are to make general moral judgments, and with good reason. Just as it is one thing to say that humans have thirty-two teeth but quite another to say that *all* humans have thirty-two teeth, it is one thing to say that lying is wrong but quite another to say that *all* lies are wrong.

[6] That is, it seems to be not only the most obvious reason for me to stop and help her (as opposed to, say, that it is in my interest to do so, or that I ought to do so), but also the most obvious explanation of why I ought to stop and help her (as opposed, to say, that my doing so would maximize utility).

[7] On this third view, Abby's being in distress is an instance of a (generic) right-making factor (*viz.*, another's being in distress), which explains why Abby's being in distress explains why I ought to help Abby. Similarly, that smoking causes lung cancer—i.e., that smoking is a lung-cancer-causing factor—explains why John's smoking—which is an instance of that (generic) factor—explains his lung cancer.

etc. After all, it is not as if it makes a difference that it is Abby who is in distress, or that it is me who is in a position to help her.[8]

To be sure, explanations in terms of general moral facts are not always satisfying. Nor are they always uncontroversial. But they are explanations nonetheless. One who accepts that harming others is wrong will (usually) be satisfied with an explanation of a particular act's wrongness that appeals to the fact that it harmed another. (In a more reflective mood, she might wonder why harming others is wrong, but that is another matter.) Of course, one who doesn't accept that, say, homosexuality is wrong won't be satisfied with an explanation of a particular act's wrongness that appeals to the supposed wrongness of homosexuality. She will insist on a different explanation, one that appeals to some principle that she does accept. (This might be either a different explanation of why the act was wrong or an explanation of why homosexuality is wrong.) But, although she can certainly deny that she was offered a *good* explanation, she cannot deny that she was offered *an* explanation.

_____

[8] This is not to say that it could not make a difference that it was someone other than Abby who was in distress or someone other than me who was in a position to help her. Suppose that a child molester, M, is in distress and that the only person in a position to help him is one of his victim's, V. It might not be the case that V ought to help M, even though one ought to help others in distress, V is such a one, and M is such an other. But even in such a case, it would not be the identity of the persons involved—V and M—but rather the *kinds* of beings they are—victim and victimizer, or molester—and the (generic) relations in which they stand to one another that would explain why it is not the case that V—whoever V may be—ought to help M—whoever M may be.

2. Moral Principles

It is often said, at least among moral realists, that it is moral principles that explain particular moral facts, such as the fact that a particular promise-keeping is obligatory. Insofar as many paradigmatic moral principles purport to state general moral facts (e.g., 'Homicide is wrong' and 'Promises should be kept'), this is hardly a problem for moral generalism. (Although, if a moral generalist wants to insist that it is moral principles as such and only moral principles as such that explain particular moral facts, she will need to identify moral principles with general moral facts, rather than with statements thereof.) But it might be thought a problem for moral generalism that the set of seemingly paradigmatic moral principles is a diverse lot. Such *moral imperatives* as 'Love your neighbor as yourself' (Leviticus 19:18, Matthew 22:39) and 'Act only in accordance with that maxim through which you can at the same time will that it become a universal law' (Kant 1785, 4: 421) might be said to be moral principles. And although such imperatives may be valid or invalid, they do not purport to state facts, let alone general moral facts. Other seemingly paradigmatic moral principles do purport to state moral facts, but not general ones. 'Everyone has the right to life, liberty and security of person' (Universal Declaration of Human Rights, Art. 3) purports to state a fact, namely, that all humans have the rights enumerated within it. So does 'An act is right if and only if it maximizes utility' (cf. Brink 1989, 237), which states that all and only maximally utile acts are right. But the facts that such *universal moral statements* purport to state are universal moral facts—moral facts about classes or their members—not general moral facts—facts about the moral attributes of kinds. Despite such diversity, moral principles are best understood as cognitive statements, true moral principles are best understood as (true)

statements of general moral fact, and moral generalists can plausibly deny that either moral imperatives or universal moral statements are moral principles.

A fundamental difficulty for the view that moral principles are imperatives is that moral imperatives are not sufficiently fundamental to be moral principles. This difficulty can be brought out by looking at a fundamental difficulty for the view of morality that it fits best, theological voluntarism. According to voluntarism, there is no fact of the matter about what is right or wrong prior to or independently of what God commands. This means that there are no moral constraints on what God could command. If God were to command us to hate rather than to love our neighbors, then it would be right to hate them and wrong to love them. On this view, God does not command what is right because it is right; rather, God commands what is right because God's commanding it makes it right. But this is implausible. It is far more plausible that God would command what is right because it is right. That God would, for example, command us to love one another because that is what we ought to do. This is a fundamental difficulty for voluntarism, and it brings out a more general difficulty for the view that moral principles are imperatives: it is far more plausible that any valid moral imperatives, whatever or whoever their source, would be derivative of true moral principles, or the facts they state, than that moral principles are themselves imperatives.[9] Thus, moral imperatives, if there are any, are not sufficiently fundamental to be moral principles.

---

[9] The same objection can be advanced against the view that moral principles are simply a subset of the rules that individuals or societies accept or follow. Such a reductive account of moral principles leaves no room for there to be a truth of the matter about what is right or wrong independently of what moral rules an individual or society accepts or follows. And it is far more plausible that valid moral rules would be grounded in true moral principles, or the facts they state, than it is that moral principles are themselves rules that individuals or societies accept or follow.

One might think that moral principles are universal moral statements, such as 'Everyone has the right to life, liberty and security of person', or 'An act is right if and only if it maximizes utility'. But there are numerous difficulties with this view, too. One is that such statements lack the modal force of moral principles. Unlike universal moral statements, moral principles imply things about the moral attributes of possible particulars, not just actual ones. (This is essential to their ability to guide conduct. It is also essential to the epistemic value of *inter alia* experience, prior cases, and thought experiments.) Another, more recalcitrant difficulty for this view is that universal statements lack the explanatory power that moral principles are supposed to have. Moral principles are supposed to explain their instances. But that everyone has the right to life, liberty, and security of person would not explain why Jane or anyone else has that right. Likewise, that an act is right if and only if it maximizes utility would not explain why Jane's or any other right (or maximally utile) act is right. The problem is that the truth of a universal statement is all but logically *irrelevant* to the truth of its instances. That all members of a certain class (e.g., the class of human beings) possess an attribute (e.g., the right to life) implies that this member of that class possesses that attribute. That much is trivial. But the former fact is utterly irrelevant to *why* this or any particular member of that class (e.g., Jane) possesses that attribute. Likewise, that two classes (e.g., the class of right acts and the class of maximally utile acts) are coextensive is utterly irrelevant to *why* any particular member of one is a member of the other. Moreover, that this member of a class has a given attribute, or that this member of one class is also a member of another, is part of what makes it true that all members of that class have that attribute, or that the

two classes are coextensive. Thus, if moral principles explain their instances, they cannot be universal moral statements.

One might try to get around the first of these difficulties—that universal moral statements lack the modal force of moral principles—by insisting that true moral principles are *necessary* truths. But whatever the merits of this response, it does nothing to address the second difficulty (which is why it is more recalcitrant than the first). Suppose that all As are Bs and that it could not have been otherwise. This fact is just as irrelevant to why this A is a B as the fact that all As are Bs. Likewise, suppose that all and only As are Bs and that it could not have been otherwise. Again, this fact is just as irrelevant to why this A is a B, or this B is an A, as the fact that all and only As are Bs. Alternately, one might object that I am misinterpreting moral principles like 'Everyone has the right to life, liberty, and security of person' and 'An act is right if and only if it maximizes utility', and that, properly interpreted, moral principles *like these* are not universal moral statements. But that is a point that I, or any generalist, could readily concede. For the immediate point is that, properly understood, moral principles are not universal moral statements. Moreover, if one denies that such principles are the kind of statement they appear to be, one owes an account of what kind of statement they are—an account on which they have the modal force and explanatory power of moral principles.

Statements of general moral fact suffer from none of these difficulties. Unlike moral imperatives, they are sufficiently fundamental to be moral principles. In fact, they seem to be just the sort of facts that the validity of a moral imperative would presuppose. The validity of 'Love your neighbor as yourself', for example, seems to presuppose that one ought to love one's neighbor as oneself, or that it is right to do so. And unlike

universal moral statements, statements of general moral fact have the modal force and explanatory power of moral principles. They have the explanatory power of moral principles because the kind of facts they state—general moral facts—explain particular moral facts. And they imply things about the moral attributes of possible particulars, not just actual ones.

It is an interesting and important question just what the modal force of a statement of general moral fact is—i.e., just what such a statement implies about the moral attributes of possible particulars. Such statements ascribe attributes to kinds. And if the attribute a true statement of general moral fact ascribes to a kind K is a necessary or essential attribute of Ks, then any K must have that attribute. Rape and genocide are plausible candidates for action kinds whose instances are wrong necessarily or essentially, but homicide is not, since some homicides seem morally permissible, and some seem morally obligatory.[10] Yet 'Homicide is wrong' seems no less plausible than 'Rape is wrong' or 'Genocide is wrong', which suggests that statements of general moral fact don't attribute necessary or essential attributes to kinds. If this is right, then the truth of 'Homicide is wrong' implies neither that all homicides are wrong nor that any homicide would be wrong. And the same is true *mutatis mutandis* for other statements of general moral fact, such as 'One ought to keep one's promises'.

Nevertheless, statements of general moral fact do have modal force. If 'Homicide is wrong' is true (that is, if homicide is indeed wrong), then an act's being a homicide is

---

[10] Would not a police officer be morally obligated to kill a would-be suicide bomber who was attempting to detonate a bomb in crowded subway car if there were no other way he could prevent the detonation?

sufficient to secure its wrongness in the absence of negating, countervailing, or overriding factors—that is, if nothing interferes. Thus, any homicide must and will be wrong unless some interfering factor, or combination of such factors, makes it permissible or obligatory. This is reflected in—and helps to explain—elements of our moral thought and practice. First, although we normally require no justification for refraining from homicide, we do normally require justification for committing homicide. Second, we don't normally require any further explanation of an act's wrongness, any explanation beyond the fact that the act is, was, or would be a homicide. Nor do we normally require any further justification for the claim that an act is, was, or would be wrong. Third, although we do think that some homicides are permissible and even obligatory, we normally think that the permissibility or obligatoriness of a particular homicide requires explanation, and that the claim that a particular homicide is, was, or would be permissible or obligatory requires justification. Fourth, when things are not as they normally are, as when an officer requires a soldier to justify his not killing an enemy soldier that he could have killed, it is because we think some negating or countervailing factor, or some combination of such factors, obtains. Speaking generally then, a statement of general moral fact, which ascribes a moral attribute to a kind, implies that particulars of that kind must and will have that attribute in the absence of interfering factors—factors that interfere with their having that attribute, either by preventing them from having that attribute or by causing them to have a different, incompatible attribute (e.g., to be right, rather than wrong).

Statements of general moral fact do have modal force, then. Indeed, one might well think that they have just the modal force that moral principles have. They also have the

explanatory power that moral principles are supposed to have, and they are sufficiently fundamental to be moral principles. Moreover, as neither moral imperatives nor universal moral statements satisfy all of these criteria, the diversity of seemingly paradigmatic moral principles is not a problem for moral generalism, or for its account of moral principles. For the moral generalist can plausibly deny that either moral imperatives or universal moral statements are moral principles, and insist that it is general moral facts that explain particular moral facts.

3. Moral Generalism and Moral Thought, Practice, and Theory

It is a virtue of moral generalism that its account of moral principles enjoys these advantages over accounts of moral principles as imperatives and universal moral statements or truths. Moral generalism also has the significant theoretical virtue of being able to explain and justify important aspects of moral thought, practice, and theory—including the various prominent and important roles that moral principles and (putative) general moral facts play therein.

We accept certain moral principles, or regard some kinds of actions (such as promise-keepings) as right and others (such as homicide) as wrong. And when we do, we take certain attitudes toward particular actions falling under these principles, or particular instances of these kinds. When, for example, we accept a principle of negative moral obligation, or regard an action kind as wrong, we normally require no justification for refraining from actions falling under it, or of that kind; rather, we require justification for committing them. We also do not normally require any explanation of why such an action is wrong beyond the mere fact that it falls under that principle, or is of that kind.

Nor do we normally require any further justification for the claim such an action is wrong. Rather, we normally think that the permissibility or obligatoriness of such an action requires explanation, and that the claim that such an action is permissible or obligatory requires justification. And when things are not as they normally are (as when an officer requires a soldier to justify his not killing an enemy soldier that he could have killed), it is because we think that some negating or countervailing factor, or some combination of such factors, obtains. We think that such actions must and will be wrong in the absence of interfering factors that render them permissible or obligatory.

Moral generalism fits these aspects of our moral thought and practice. It also explains and justifies them, because it holds that general moral facts (e.g., that homicide is wrong) explain particular moral facts (e.g., that this homicide is wrong), and because general moral facts have just these implications. If a kind K has an attribute A, then any particular K must and will have A in the absence of interfering factors—factors that interfere with its having A, either by preventing it from having A or by causing it to have a different, incompatible attribute. Thus, if homicide is indeed wrong, then any particular homicide must and will be wrong unless some interfering factor, or combination of such factors, makes it permissible or obligatory.

Moral generalism, including its general moral fact account of moral principles, also fits, explains, and justifies the prominent and important roles that appeals to moral principles and (putative) general moral facts—collectively *general moral considerations*—play in diverse areas of moral thought and practice. To be sure, much moral and other practical thought is "intuitive", or at least not consciously deliberative or inferential. Nevertheless, we appeal to general moral considerations in various areas of moral thought

and practice. We appeal to them in, among other things, moral deliberation and advising, moral justification and criticism, and moral persuasion and argument. They also play prominent and important roles in moral learning and education. And we regard them as having normative force; in some cases, we even strive to live by them. Moral generalism fits these aspects of our moral thought and practice. What is more, it explains and justifies them.

It is common for individuals engaged in practical deliberation to reflect on general moral considerations. This is most obvious, and perhaps most common, when we are faced with what are obviously moral choices, as was the pupil of Sartre's who sought his advice about whether to leave for England to join the Free French Forces or to remain in France to help his mother "to live" (Sartre 1948, 33). Although Sartre presents the young man's choice as one "between two kinds of morality," one "of sympathy, of personal devotion" and one "of wider scope" (33), it is much more naturally seen as a choice between fulfilling one obligation and fulfilling another, conflicting obligation. He ought to leave and join the Free French, but he ought also to stay and help his mother. Thus, the case is much more naturally described as one in which one must decide which of two moral considerations outweighs the other in one's particular circumstances. And the pupil's choice is a difficult one for the very reason that, on reflection, he thinks that conflicting moral considerations bear on it but is not sure which outweighs the other. These moral considerations are *general* ones because the pupil bears the obligations he does *qua* instance of one or more morally significant or relevant kinds (*qua* Frenchman and son, perhaps), rather than *qua* the particular person he is—as is shown by the fact that we do not need to know *who* he is to understand his predicament.

Less obviously, and perhaps less commonly, we reflect on general moral considerations when faced with less-obviously moral choices. Suppose Debbie is deliberating about how to spend her holiday bonus. Although this is not obviously a moral choice, Debbie might think to herself "I really ought to help those who are less fortunate than I am," which invokes the principle, 'One ought to help those less fortunate than oneself'. Moreover, reflection on general moral considerations might lead her to donate her bonus to a local charity, rather than to spend it on herself. We also reflect on general moral considerations in deciding whether or not a choice is a moral one. Whether we conclude that it is such a choice depends on whether we conclude that any such considerations bear on it. Perhaps Debbie did not see the choice of how to spend her holiday bonus as a moral one until she realized that a moral principle bore on it. Perhaps she saw it as a merely prudential choice.

Even more commonly than reflection on general moral considerations forms part of individual deliberation, it forms part of collective deliberation and advising. Suppose, for example, that a government commission is formed and charged with making policy recommendations regarding some particularly controversial group of medical research programs, say those involving the use of (human) embryonic stem cells. It would hardly be surprising to hear Roman Catholic bishops urge the commission to consider the "sanctity of human life" in their deliberations. Nor would it be uncommon for such a body to give this and other general moral considerations significant weight in its deliberations. Nor, again, is such reflection limited to cases that present obviously moral choices. Many public policy choices have moral implications that are not obvious, for example.

If reflection upon general moral considerations is common in moral deliberation and advising, it is even more common to appeal to such considerations in justifying and criticizing actions. Suppose that Debbie donates her holiday bonus to a local charity, and that her friend Drew asks her why she did so. She may well say that it was because one ought to help those who aren't as fortunate as oneself. Moreover, Drew may well agree with her, and judge that she acted rightly. Or suppose that our government commission recommended federal funding for embryonic stem cell research. It might justify its recommendation by citing the great promise such research holds for relieving human suffering, thus implying that it is right to relieve human suffering. Our Roman Catholic bishops, on the other hand, might criticize the commission's recommendation for failing to accord the "sanctity of human life" the respect that it deserves.[11]

Such appeals to general moral considerations in justifying and criticizing actions are common, whether those actions are our own or those of other people or institutions (e.g., governments and businesses). In a similar fashion, appeals to general moral considerations are common in moral persuasion and argument. Debbie might try to persuade Drew that he too should donate his holiday bonus to charity, citing our obligation to help those less fortunate than ourselves. Or our bishops, unhappy with the outcome of the commission's deliberations, might argue that the commission should reconsider its recommendations, or that policy makers should not follow them, citing our

---

[11] Compare "Criticism comes from both sides", *Arizona Republic*, 10 August 2001, quoting Joseph A. Fiorenza, president of the U.S. Conference of Catholic Bishops, as calling President Bush's decision to allow federal funding of research involving existing stem cell lines derived from already-destroyed human embryos "morally unacceptable". "We hope", the *Republic* quotes Fiorenza as saying, "that President Bush will return to a principled stand against treating some human lives as nothing more than objects to be manipulated and destroyed for research."

obligation to respect the sanctity of human life. Moreover, these are just some of the ways in which general moral considerations and appeals thereto figure in moral thought and practice.

General moral considerations also have a prominent and important place in moral education, which is commonly thought to involve the teaching and inculcation of moral principles. We teach our children to share, to be honest, to help those who are less fortunate than we are. And we teach them not to lie, not to cheat, and not to steal. Presumably, we are teaching them that sharing, being honest, and helping others are right, and that lying, cheating, and stealing are wrong. And we appeal—explicitly and implicitly—to such general moral considerations when we explain to them why they ought or ought not do this or that. "Why can't I take the doggie home, daddy?", Emma might ask. "Because, Emma, the doggie belongs to Kiera." "So?", Emma might wonder. "It's not right to take something that belongs to someone else, Emma. It's stealing, and stealing is wrong." Similarly, when we seek to educate ourselves morally, we seek to discover true moral principles and to live by them. Gandhi and Martin Luther King, Jr. embraced principles of pacifism and non-violence, and they lived by those principles, even when actively resisting bigotry, hatred, and violence.

All of these practices are closely related, of course. All involve reflection on or appeal to reasons for action, specifically normative reasons.[12] Practical deliberation

---

[12] The contrast I have in mind here is between normative, or justifying, reasons and (merely) explanatory reasons. While elements of a person's upbringing might be (merely) explanatory reasons for why he is, say, a serial killer, they would not provide normative or justifying reasons for his being one. One would not say that someone was *justified* in being a serial killer because he was sexually abused as a child.

involves reflecting on the reasons one has for acting one way or the other. Practical advice-giving, justification, and argument all involve offering reasons for action. And practical persuasion often involves the offering of such reasons.[13] In moral deliberation and advising, in justifying and criticizing actions and moral judgments, and in moral persuasion and argument, it is common to treat general moral considerations as having normative import. In a similar fashion, our moral education practices treat such considerations as having normative import, as being relevant to how we should act in particular cases, and to how we should live our lives. Some among us even strive to live by moral principles.

The generalist account of these aspects of moral thought and practice, and of the roles and normative import they accord general moral considerations, is that they reflect the relations that obtain between and among moral principles, general moral facts, and particular moral facts, together with the normative character of these. True moral principles state or express general moral facts, including the facts about the rightness and wrongness of action kinds, which in turn explain particular moral facts, including the rightness and wrongness of particular actions. And inasmuch as the rightness, wrongness, and other moral properties of actions are normative or normatively relevant ones (i.e., relevant to what we have reason to do), it only makes sense to reflect on and appeal to general moral considerations—moral principles and (putative) general moral facts—when deliberating and advising, when justifying and criticizing actions, when

---

[13] It also frequently involves a number of other things—hence, the distinction between argument and rhetoric.

attempting to persuade and when arguing with others, when educating oneself and others, and when trying to live as one should. Thus, moral generalism, including its account of moral principles, not only fits these aspects of moral thought and practice; it also makes sense of them. It explains and justifies them.

Moral generalism's ability to explain and justify these aspects of our moral thought and practice is a considerable virtue. But its greatest virtue may be its ability to explain and justify the central role that moral principles play in moral explanation—in explaining moral phenomena such as the rightness of particular actions, the justness of particular social systems and institutions, and the nature and extent of moral rights. We typically explain why particular right actions are right—and why particular wrong actions are wrong—by subsuming them under *principles of obligation*. We typically explain why particular social systems and institutions are just or unjust by subsuming them under *principles of justice* for such systems and institutions. More generally, we typically explain particular moral facts with moral principles—though this is sometimes implicit, as when one cites a particular fact (e.g., that Abby is injured) to explain a particular moral fact (e.g., that I ought to stop and help Abby). Moral principles play this explanatory role both in commonsense moral thought and practice and in moral theory, but it is arguably their primary moral-theoretic role.[14] A moral theorist might—at some cost—claim that the principles of her theory need not play all of the roles that moral principles play, but her theory must include principles that explain the rightness of particular right acts—at least

---

[14] Indeed, some opponents of the kind of radical particularism that rejects moral principles altogether reject such particularism on the grounds that moral principles are necessary to explain particular moral facts (Sidgwick 1907, 98-103; Brink 1994a, 183, 188-90).

so long as she claims that her theory aims to explain what makes right acts right, which many contemporary moral theorists do.[15] It is thus a considerable virtue of moral generalism that it offers a ready and plausible explanation of this practice—the practice of appealing to moral principles to explain moral phenomena. According to moral generalism, general moral facts explain particular moral facts and moral principles state (putative) general moral facts. Thus, we can explain why particulars bear the moral attributes they do by subsuming them under moral principles.

4. From Moral Generalism to Moral Dispositionalism

But what sort of facts are general moral facts, such that they are sufficiently fundamental and have the modal force and explanatory power of moral principles? I have said that they are facts about the moral attributes of kinds; but what sort of attributes are we talking about? They must be the right sort of attributes if general moral facts are to be what moral principles state. But which sort of attributes are those? And which kinds bear which moral attributes? This last question is one for normative moral theory, and thus beyond the scope of this inquiry. But the former is a question for moral metaphysics, or the metaphysics of morality. And it is one this dissertation aims to answer.

---

[15] Thus, for example, if an act-utilitarian claims that act-utilitarianism is a theory of what makes right acts right and not a model for moral reasoning or a decision-making procedure, her "principle of utility" will need to explain why right acts are right, even though it need not be of much, or any, use to actual agents trying to decide what to do (cf. Bales 1971; Brink 1989, 256-62).

The answer to be given is that they are dispositions, that general moral facts are facts about the dispositional attributes—including the moral tendencies and powers—of kinds of things—including kinds of beings (e.g., agents and persons), actions (e.g., homicides and promises), and states of affairs (e.g., that one promised to ϕ and that someone is injured). This answer is suggested by Ross, who appeals to dispositions (specifically, tendencies) in order to explain the nature of moral principles and conflicts between them. In the next chapter, I offer an account of the metaphysics of conflicts of obligation, which is something any adequate account of the metaphysics of morality must do, and I lay the groundwork both for a fuller understanding of moral holism and for the development of a moral dispositionalism that explains both moral holism and moral generalism.

CHAPTER 2

CONFLICTS OF OBLIGATION

Sometimes our obligations conflict, as the following case illustrates:

> I promise to meet Andy at the pub at six, after I return from my hike and
> have a chance to clean up. But during my hike, I come upon Abby, who
> is injured and needs help. If I stop to help her—to see to her injury,
> accompany her to the trailhead, and see that she receives proper medical
> attention—I won't make it to the pub in time to meet Andy for a pint,
> nor will I be able to let him know this.

I ought to keep my promises, and I ought to help those in distress. I have two

obligations: one to meet Andy and one to help Abby. Normally this would be no

problem, but given the circumstances, I cannot fulfill both of these obligations. They

conflict.

Now I take it that this particular *conflict of obligations* is easily resolved. I ought to help

Abby *rather* than meet Andy, even though I will break my promise to Andy if I do so. In

this case, my obligation to help those in distress outweighs my obligation to keep my

promises. Such conflicts occur every day, and although all too many are not so easily

resolved (think of Sartre's pupil), there is nothing mysterious about them.[1] There is

---

[1] Nevertheless, some would deny such conflicts are possible. Kant thought conflicts of
obligation were literally inconceivable (Kant 1797a, 6: 224), and Hare seems to agree (Hare 1981,
26). Prichard evidently changed his mind more than once on the matter (Prichard 1912, 14 n. 7;
Prichard 2002, 78-79, 287). And despite offering us the resources to construct an account of the
metaphysics of such conflicts (see below), Ross, too, may have thought them impossible (Ross
1939, 84-86). In any case, there is nothing in the concept of an obligation that prevents such
conflicts. Otherwise, I could not have conflicting social, legal, or contractual obligations, all of

simply a difference between something one ought to do and what one ought to do *rather than anything else*. And while I have both an obligation to help Abby and an obligation to meet Andy, helping her is what I ought to do rather than anything else. It is my *obligation overall*.[2]

Although there is nothing mysterious about conflicts of obligation, they do raise significant metaphysical questions. If I ought to help Abby rather than meet Andy, then I ought *not* meet Andy. But, given that I have an obligation to meet Andy, it follows that my meeting him would be both right and wrong. This raises three questions. First, how can my meeting Andy have contrary moral properties? How can it be both right and wrong?[3] Second, in virtue of what is my meeting Andy wrong? What makes it wrong? Third, in virtue of what is my helping Abby what I ought to do rather than anything else? What makes helping her my obligation overall? If commonsense morality is (or even just might be) correct in recognizing an irreducible plurality of moral obligations that conflict in particular cases, then we should to be able to answer these questions. Moreover, those moral pluralists who believe that our obligations can and do conflict in particular cases

_____

which I can have. Nor could such obligations conflict with one another, or with my moral obligations, which they can.

[2] Foot distinguishes between "type 1 ought statements" and "type 2 ought statements" and identifies conflicts of obligation as involving conflicts between two or more of the former, each of which is true (Foot 1983, 255-56). But she doesn't offer an account of the underlying metaphysics of such conflicts, as I do here.

[3] This question is also raised by cases in which a particular action is an instance of two distinct morally significant or relevant kinds one of which is right and the other of which is wrong, as when a particular action is both an instance of lying (say, to a man intent on murder) and an instance of protecting the innocent from harm (cf. Kant 1797b, 8: 423). One might think the answer is obvious in this case: that an action can be wrong in one respect but right in another (see below).

need an account of the metaphysics of conflicts of obligation in order to answer the objection lodged by some monists that their theories are self-contradictory (Ross 1939, 83-86).[4]

Ross offers a widely accepted account of the metaphysics of conflicts of obligation (Ross 1939, 83-86; Ross 1930, 17-20). Unfortunately, it appears to be an eliminativist one, for Ross appears to hold that conflicts of obligation are something else entirely and that genuine conflicts of obligation are impossible. On Ross's account, conflicts of obligation are conflicts between what he calls "*prima facie* obligations", rather than conflicts between what he calls "obligations", and conflicts between the latter are impossible. Moreover, although Ross himself appeals to neither notion, it is standard to gloss *prima facie* obligations as, or in terms of, moral reasons or right-making factors; and neither conflicts between moral reasons nor conflicts between right-making factors are obviously conflicts between obligations. Indeed, they might well be taken to be conflicts between the grounds of obligations, rather than between obligations themselves (cf. Kant 1797a, 6: 224). What is more, Ross's account of is generally taken to be atomistic (e.g., Brink 1994b, 219), which would make it inconsistent with moral holism.

---

[4] *Moral pluralists* believe there is an irreducible plurality of "fundamental and underivative" (Ross 1930, 57) obligations or moral principles, while *moral monists* believe there is a single such obligation or principle. It is not obvious that moral monists do not also need an account of the metaphysics of conflicts of obligation; for even putting to one side the question of conflicts between derivative obligations, there is the possibility that a single, fundamental obligation might conflict with itself. Suppose that my only obligation is to promote the good. Might I not then have several obligations to promote several goods, not all of which I can promote? Or suppose that my only obligation is to promote happiness. Might I not then have several obligations to promote the happiness of several persons or other sentient beings, not all of which I can promote? It is certainly the case that, as Foot notes, one can be obligated to keep each of two promises (but not both of them) even though it is impossible for one to keep them both (Foot 1983, 254).

Appearances and standard glosses notwithstanding, Ross offers us the resources to construct an intuitive and plausible account of the metaphysics of (genuine) conflicts of obligation that is perfectly consistent with moral holism. In this chapter, I look to some unappreciated aspects of Ross's moral metaphysics to explicate this account, along with what I take to be the correct understandings of *prima facie* obligations and of what it is for a particular action to be *prima facie* obligatory—i.e., of *prima facie* rightness or obligatoriness. In particular, I look to the roles that degrees of rightness and the dispositional notion of a tendency play in Ross's moral metaphysic. On this account, *prima facie* obligations are obligations under a different, overtly dispositional description,[5] while *prima facie* obligatoriness is the tendency to be obligatory overall, or to be the most right course of action open to an agent in a situation. Thus, conflicts of obligation are conflicts between genuine obligations. Moreover, as we will see in Chapter 5, it is a short step from this dispositional account of *prima facie* obligations to moral dispositionalism, which not only accommodates both moral generalism and moral holism but also explains them both.

---

[5] In taking *prima facie* obligations to be obligations, my understanding of *prima facie* obligations is closer to those suggested by Audi (1996, 103-4), Shope (1965, 280 & n. 1), and Prichard (2002, 287), than it is to those suggested by many others (although neither regard them as obligations under an overtly dispositional description). For not only do I eschew attempts to understand *prima facie* obligations as, or in terms of moral reasons, I also eschew attempts to understand them as, or in terms of, right-making factors or properties. Cf. Bakhurst 2000, 163; Brink 1994b, 216; Crisp 2000, 46-47 & n. 46; Dancy 1993, 97; Dancy 2000, 131; Dancy 2004, 18; Ewing 1959, 109-11; Frankena 1973, 26; Gaut 2002, 138; Hooker 2000; Hooker 2002, 178-83; McNaughton 1996, 78-79. These three concepts are certainly interrelated, but there are different ways of understanding the relations between *prima facie* obligations and moral reasons, and if *prima facie* obligations just are obligations, then a right-making factor is one that makes an action (*prima facie*) obligatory, not a *prima facie* obligation.

1. <u>Fixing Ideas</u>

English does not supply distinct terms for what I called an obligation and an obligation overall, for something one ought to do and what one ought to do rather than anything else. In Ross's terminology, the obligation to keep one's promises is a "*prima facie* duty" or "*prima facie* obligation", as is the obligation to help those in distress (Ross 1930, 19-20; Ross 1939, 84-86).[6] I can also be said to have a *prima facie* obligation to meet Andy and a *prima facie* obligation to help Abby, and both my meeting Andy and my helping Abby are "*prima facie* right" or "*prima facie* obligatory". (Only very rarely does Ross speak of a particular action as being a *prima facie* duty or a *prima facie* obligation, though this usage is now fairly common.) Had circumstances not conspired to make it impossible for me to fulfill both my *prima facie* obligation to meet Andy and my *prima facie* obligation to help Abby, I would have been obligated overall to perform both of these (*prima facie* obligatory) actions. (In Ross's terminology, I would have been "obliged" or "obligated" to perform both of them.) But in the case I described, meeting Andy was not my obligation overall; rather, helping Abby was my obligation overall. As well as being *prima facie* obligatory, my helping Abby was obligatory overall.

---

[6] Ross uses the terms '*prima facie* duty', '*prima facie* obligation', 'claim' (which is Prichard's term), and 'responsibility' (which is E.F. Carritt's term) interchangeably (Ross 1939, 43-44, 84-85, 315.) Ross's use of 'right' and 'rightness' is also liable to confuse. In *The Right and the Good*, Ross announces he will ignore a slight difference in meaning between the words 'right' and 'obligatory' and use the former for the latter (1930, 4). But in *Foundations of Ethics* (1939), he makes no such stipulation and uses the words differently. Thus, in the former, 'right' means 'obligatory', 'rightness' means 'obligatoriness' or 'obligation' (depending on the context), and so on. Ross also prefers the word 'duty' in the former work but 'obligation' in the latter. Neither he nor I distinguish between duties and obligations.

The case of Abby and Andy is one in which one *prima facie* obligation outweighs

another, but there are other kinds of cases. An instance of the same general obligation

can outweigh another instance of the same obligation, for instance—as when one's

obligation to keep one promise outweighs one's obligation to keep another. Two equally

weighty obligations can also conflict, as can two equally weighty instances of the same

obligation. The following case illustrates the latter possibility:

> A pair of twin boys is born, each suffers from a physical defect that will
> end his life in short order unless an operation is performed immediately.
> Only one surgeon who could perform the operations is at hand, and the
> circumstances are such that she can only save one of the boys. The other
> will die. The father is absent, the mother unconscious. The surgeon must
> make the decision about which boy's life to save. There is no medical
> reason for saving one rather than the other, as neither boy has a better
> prognosis, neither surgery will be easier than the other, neither surgery will
> require fewer resources that might be used to help others, and so on.[7]

When, from the standpoint of moral deliberation, an agent asks, 'What is *the* right thing

for me to do?', she is asking, 'What am I obligated overall to do?' In the case of Abby

and Andy, my obligation overall was to help Abby, inasmuch as my helping Abby was the

one action I ought to perform rather than any other. But here, there is no *one* action the

surgeon ought to perform rather than any other. She has a *prima facie* obligation to

---

[7] Some would regard these further considerations as moral ones, and so relevant to what the
surgeon is obligated to do. Others would regard them as non-moral considerations, and so
irrelevant to what the surgeon is obligated to do. But one who regards them as morally irrelevant
need not regard them as ethically irrelevant. One could say that if these further considerations
favored saving one boy rather than the other, it would be better—and hence more
reasonable—for the surgeon to save that boy, even though she would not be obligated to save
him rather than the other boy. As Williams points out, "someone who was not a utilitarian . . .
might think that utility indeed was the only possible universal resolver of conflicts [of obligations
and of other values]" (1981, 78). Of course, as Williams notes (73-74), to resolve a conflict of
obligations is not to eliminate or annul one of those obligations or to show it to be merely
apparent, rather than real.

operate on the one boy and a *prima facie* obligation to operate on the other boy; but neither operation is obligatory overall, because neither is such that she ought to perform it rather than the other. Rather, her obligation overall is to perform one or the other of the two operations.[8]

Such cases show that we need to distinguish—as Ross does—between an *action* an agent ought to perform rather than any other and *what* an agent ought to do rather than anything else, for the surgeon's obligation overall is not a particular action, nor is it to perform a particular action. If we just look at cases like that of Abby and Andy, it would seem that we could say that when an agent's *prima facie* obligations conflict, her obligation overall is the weightiest of her *prima facie* obligations (cf. McNaughton 1988, 198; Prichard 2002, 287). But the surgeon's obligation to operate on one or the other of the two boys is not an obligation to fulfill the weightiest of her *prima facie* obligations. It is, rather, an independent consequence of her not being able to fulfill both of two maximally weighty *prima facie* obligations. It is an obligation to pursue a *course of action* that is not identical to fulfilling one of her *prima facie* obligations (although *in this case* she would fulfill this obligation by fulfilling one of her *prima facie* obligations.)

It might be suggested that an agent's obligation overall is to fulfill the weightiest of her *prima facie* obligations *or* one or the other of the weightiest of them, if no one of them is weightier than all the others. But there may be cases in which the right course of action involves a compromise between the demands of conflicting *prima facie* obligations.

---

[8] Here Ross would say that either operation would be "right", but that neither is "obligatory". And he would say that her "obligation" is to perform one or the other of these *prima facie* obligatory actions—i.e., to operate on one or the other of the two boys (Ross 1930, 3-4; Ross 1939, 43-44).

Suppose that I owe two creditors $100 each, but have only $100. My obligation overall

may be to pay each of them $50, or to pay one $40 and the other $60, rather than to pay

one of them $100. If so, it is neither to fulfill the weightiest of my *prima facie* obligations

nor to fulfill one or the other of my weightiest *prima facie* obligations. It is simply the

most right thing I can do under the circumstances (Ross 1939, 84). And that is what our

obligation overall is in any case: the most right thing we can do given our *prima facie*

obligations and the circumstances, the most right course of action open to us. Sometimes

the most right course of action is a particular action, but sometimes it is to perform one

or the other of a number of maximally right actions. And even when it is a particular

action, that action may not fulfill one of our *prima facie* obligations.

2. Conditional Obligatoriness

Now that we have a fix on the relevant ideas, we can fruitfully consider the nature

of *prima facie* obligations, and of *prima facie* obligatoriness. The obvious place to begin is

with what Ross says about the term '*prima facie* duty':

> I suggest '*prima facie* duty' or 'conditional duty' as a brief way of referring
> to the characteristic (quite distinct from being a duty proper) which an act
> has, in virtue of being of a certain kind (e.g. the keeping of a promise), of
> being an act which would be a duty proper if it were not at the same time
> of another kind which is morally significant. (Ross 1930, 19)

This is Ross at his least perspicuous, and it is easy to misunderstand both what Ross says

here and why he says it (i.e., what Ross is doing in this passage). Dancy, for example,

takes this to be Ross's "official definition of a prima facie duty" (Dancy 1993, 97)—that

is, his official account of what a *prima facie* obligation *is* exactly. He also takes it to be an

attempt to define what he calls a contributory reason (Dancy 2004, 18). But it is neither.

It is, rather, a suggestion for how we might refer to a *characteristic* or property possessed by some particular acts.[9] Elsewhere, Ross refers to this same property as 'conditional or *prima facie* rightness', and describes it as:

> the property, not necessarily of being right but of being something that is right if the act [that bears it] has no other morally significant characteristic (such as that of causing much pain to another person). (Ross 1930, 138)

By 'right' here, Ross means obligatory overall (see note 6), and in his later work, Ross prefers the term '*prima facie* obligatoriness' (Ross 1939, 85-86).

We can break down what Ross says about *prima facie* obligatoriness in these passages as follows. First, he says that the property itself is that of being obligatory overall "where other considerations do not enter into the case" (1930, 135). That's the basic characterization. Second, particular actions bear this property in virtue of being instances of morally significant or relevant kinds (promise-keepings, helpings of others in distress, etc.) (Ross 1930, 19, 29; Ross 1939, 90). Thus, F-actions are *prima facie* obligatory just in case F-actions are, *in virtue of being F*, obligatory overall where other considerations do not enter into the case. Third, a particular F-action is *prima facie* obligatory just in case it is such as to be obligatory overall if morally relevant considerations other than its being F do not enter the case, which would be for it to also be G in virtue of being an instance of a second morally significant kind—namely, G-actions.

Not only is this a rather cumbersome way of characterizing *prima facie* obligatoriness, it is not at all a good characterization of it. The first problem is that it obscures three facts that ought to be obvious, and that Ross himself rightly assumes throughout his

---

[9] In the case of possible actions, it is a property they would have if performed. For convenience, I ignore this complication.

work. The first of these is that particular *prima facie* obligatory actions are not *just* obligatory overall where other considerations do not enter into the case. They are also obligatory overall *unless* something prevents them from being so (cf. Ross 1930, 138). On Ross's own view, any particular *prima facie* obligatory action is, and indeed must be, obligatory overall unless something prevents it from being so.

The second fact this characterization of *prima facie* obligatoriness obscures is that the mere presence of other considerations in a case—the mere fact that other considerations do enter the case—is not sufficient to prevent a *prima facie* obligatory action from being obligatory overall. Some considerations add to the stringency of *prima facie* obligations in particular cases. A *prima facie* obligation to keep one's promise to visit a gravely ill friend is (at least *ceteris paribus*) more stringent, and thus less easily overridden by another *prima facie* obligation, than is a *prima facie* obligation to keep one's promise to go to the theater with friends (Ross 1939, 100). But even a consideration that reduces the stringency of a *prima facie* obligation to perform an action does not necessarily prevent that action from being obligatory overall. Perhaps the *prima facie* obligation to speak the truth to others is less stringent when the other in question is a habitual liar than when he is not, but it would not follow that it cannot be obligatory overall to fulfill this obligation when the other in question is such a liar (cf. Ross 1930, 54-55). Moreover, whether a particular *prima facie* obligatory action is obligatory overall ultimately depends in part on the *comparative* stringency of the agent's *prima facie* obligation to perform it (Ross 1930, 41;

Ross 1939, 85).[10] Thus, even if considerations reduce the stringency of an agent's *prima facie* obligation to perform a particular action, it still may be the most stringent, or the weightiest, of her *prima facie* obligations.

The third fact this characterization of *prima facie* obligatoriness obscures is that there may be different kinds of morally relevant considerations, or morally significant action kinds. It may be tempting to say that "what we mean by a 'prima facie' obligation . . . is what would be an over-all obligation if there were no conflicting stronger obligations" (Brandt 1959, 368).[11] But consider the fact that *prima facie* obligatory actions can be prevented from being obligatory overall by being physically impossible to perform. If a book I borrowed from you is destroyed in a fire, I cannot be obligated overall to return it to you (though I may be obligated overall to replace it or compensate you for the loss).[12] But whereas actions may be *prima facie* obligatory in virtue of being promise-keepings, they are surely neither *prima facie* obligatory nor *prima facie* wrong in virtue of being physical impossibilities. Physically impossible actions may be a morally significant kind of action, but they are a morally significant kind of a different sort than promise-keepings.

---

[10] Ross writes that "whether an act is a duty proper or actual duty depends on *all* the morally significant kinds it is an instance of," and that "*being* one's duty is a toti-resultant attribute, one which belongs to an act in virtue of its whole nature and nothing less than this" (Ross 1930, 19-20, 28). Given that whether an action is obligatory overall ultimately depends in part on the *comparative* stringency of the agent's *prima facie* obligation to perform it, I suggest that what he means is that an action's overall moral status—its overall rightness or wrongness—depends on all its properties, including its relational properties (cf. Prichard 2002, 286).

[11] Cf. Brink 1994b, 216: "an all-things-considered obligation is an undefeated prima facie obligation."

[12] Ross claims that in such a case I would not be even *prima facie* obligated to return it to you (Ross 1939, 109).

Given all this, it would be better to say that for a particular action to be *prima facie* obligatory is for it to be such as to be obligatory overall if nothing prevents it from being so. Unlike what Ross says, this captures more of what it is to be *prima facie* obligatory and clearly places no limits on what sorts of considerations can prevent a *prima facie* obligatory action from being obligatory overall, or what sorts of morally significant action kinds there are other than *prima facie* obligatory ones. But though this is a better characterization of *prima facie* obligatoriness than the one Ross gives, it still isn't an adequate one.

3. <u>The Tendency to be Obligatory Overall</u>

It isn't adequate because it suffers from the second problem that Ross's own characterization of *prima facie* obligatoriness suffers from. It doesn't capture the fact that when a particular *prima facie* obligatory action fails to be obligatory overall, the *prima facie* obligation to perform it still makes a claim on the agent to fulfill it. To put it another way, the actions we have *prima facie* obligations to perform are not *merely* such as to be obligatory overall if nothing prevents them from being so; rather, they are such that we ought to do them. This is an essential feature of *prima facie* obligations, and neither this nor Ross's own characterization of *prima facie* obligatoriness captures it.[13]

---

[13] Thus, they both invite a misunderstanding that Ross is at pains to avoid: that *prima facie* obligations are merely apparent obligations (Ross 1930, 20; Ross 1939, 85; cf. Hare 1981, 38-45; Foot 1983, 257-58). The distinction between a *prima facie* obligation and an obligation is not the distinction between an apparent obligation and a real obligation. The latter is a perfectly good distinction, but it is applicable to *prima facie* obligations, as well as to other obligations. An action may appear to be *prima facie* obligatory and yet turn out not to be.

Suppose that my meeting Andy is merely such as to be obligatory overall if nothing prevents it from being so. Then when my meeting Andy is prevented from being obligatory overall, the *prima facie* obligation to keep my promises might no longer make any claim on me to meet him. But that isn't what happens. My *prima facie* obligation to meet Andy is not cancelled or nullified by my not being able to both meet him and help Abby. Nor is it transformed into a different obligation (cf. McNaughton 1988, 198-99). Rather, it is *outweighed* by my *prima facie* obligation to help Abby. Similarly, the surgeon's *prima facie* obligation to help the one boy is not nullified, either by her *prima facie* obligation to help the other boy or by her obligation overall to help one or the other. In each case, each of the relevant *prima facie* obligations makes a claim on the agent whose *prima facie* obligation it is—a claim to perform an action that will fulfill it. The mere fact that each of those actions is such as to be obligatory overall if nothing prevents it from being so is consistent with things being otherwise, and thus a characterization of *prima facie* obligatoriness in terms of that fact alone is inadequate. We need a characterization of *prima facie* obligatoriness that captures this essential feature of *prima facie* obligations.

As I said, Ross's characterization of *prima facie* obligatoriness has this same defect. And this seems to be the basis for both Williams's complaint that Ross's account of moral conflict "eliminate[s] from the scene the *ought* that is not acted upon" (Williams 1965, 175)[14] and Dancy's complaint that Ross "tells us only what effect 'being a prima facie duty in virtue of property F' has when F is the only morally relevant property" and

---

[14] Williams supposes that the claim that an agent's *prima facie* obligation makes on her "must surely be a claim for consideration as *the only thing that matters*, a duty" (my emphasis) (Williams 1965, 176). But while *prima facie* obligations may make claims to be duties, they do not make claims to be the only thing that matters.

"has nothing to say at all about what it is like to 'be a prima facie duty in virtue of

property F' in any other situation" (Dancy 1993, 97). But Ross both recognizes this

defect and addresses it:

> It is necessary to say something by way of clearing up the relation between *prima facie* duties and the actual or absolute duty to do one particular act in particular circumstances. If, as almost all moralists except Kant are agreed, and as most plain men think, it is sometimes right to tell a lie or to break a promise, it must be maintained that there is a difference between *prima facie* duty and actual or absolute duty. When we think ourselves justified in breaking, and indeed morally obliged to break, a promise in order to relieve someone's distress, we do not for a moment cease to recognize a *prima facie* duty to keep our promise, and this leads us to feel, not indeed shame or repentance, but certainly compunction, for behaving as we do; we recognize, further, that it is our duty to make up somehow to the promisee for the breaking of the promise. We have to distinguish between the characteristic of being our duty and tending to be our duty. Any act that we do contains various elements in virtue of which it falls under various categories. In virtue of being the breaking of a promise, for instance, it tends to be wrong; in virtue of being an instance of relieving distress, it tends to be right. (Ross 1930, 28-29)

Here Ross clarifies his view: *prima facie* obligatory actions are not merely such as to be

obligatory overall where other considerations do not enter into the case, nor are they

merely such as to be obligatory overall if nothing prevents them from being so; they also

tend to be obligatory overall. Moreover, their so tending is what accounts for the fact that

outweighed *prima facie* obligations still make claims on us, and that failing to fulfill those

claims can give rise to residual obligations—i.e., obligate us to do other things.

Ross's claims about particular actions tending to be right and tending to be wrong

have been read in different ways, as has his claim that *prima facie* obligations are things

that tend to be our duty. Strawson wrongly assumes that they must be read statistically,

or as claims about frequencies. He wrongly assumes that the only thing the claim that

promise-keepings tend to be right could mean is that most promise-keepings are right,

and he wrongly assumes that the claim that a particular promise-keeping tends to be right would have to be the nonsensical claim that the particular action in question is usually right (Strawson 1949, 29). A quick glance at the Oxford English Dictionary suffices to show that these claims admit of different readings. The first potentially relevant definition it gives for 'tend' is "to have a motion or disposition to move towards, and derived senses."[15] Moreover, Ross's claims won't bear the readings Strawson tries to force on them.

In response to Strawson, A.C. Ewing says that "the notion of tending is . . . the notion that, other things being equal, A will be B, though this may be prevented by counteracting circumstances, the tendency being regarded as based on something positive in A" (Ewing 1959, 109). But this isn't right either, at least not if what Ewing is after is Ross's notion of tending. On Ewing's *ceteris-paribus* reading, Ross's claim that a particular promise-keeping tends to be obligatory overall adds nothing to the claim that it is such as to be obligatory overall unless it is prevented from being so. Thus, Ewing's notion of tending can't be Ross's, because it can't do the work that Ross's does. For much the same reason, the "extended analogy" that Paul Pietroski draws between "statements ascribing prima facie obligations to agents and nonmoral statements which are true only ceteris paribus" (Pietroski 1993, 490; cf. Brink 1994b, 216-20) is too weak an analogy. For what is true only *ceteris paribus* is not true when *ceteris* isn't *paribus*, which in the present case is when one *prima facie* obligation outweighs another.

---

[15] The first definition it gives for 'tendency' is "the fact or quality of tending to something; a constant disposition to move or act in some direction or toward some point, end, or purpose; leaning, inclination, bias, or bent toward some object, effect, or result." It gives *no* definition for either word corresponding to Strawson's usage.

On Ross's view, the tendencies of particular *prima facie* obligatory actions to be obligatory overall are what account for the fact that outweighed *prima facie* obligations still make claims on us, and thus for the fact that failing to fulfill such an obligation can give rise to residual obligations. Thus, Ross's tendencies must be dispositions—not to move toward something, but to be obligatory overall.[16] And to tend to be a duty must be to be disposed to be what one ought to do rather than anything else, or to be obligatory overall.[17]

To understand Ross's view, it helps to know how something can fail to manifest a disposition normal for things of its kind—a type-normal disposition. This can happen in one of two ways. First, some background condition could prevent it from manifesting that disposition, as the absence of oxygen prevents a match from lighting when struck without affecting its flammability. Second, some background condition could prevent it from possessing that disposition, as being wet prevents a match from being flammable

---

[16] In fairness to Pietroski, he does say that "just because *M*'s prima facie obligation to do *a* does not *manifest* [my emphasis] itself in its being the case that *M* ought$_{act}$ [i.e., ought all-things-considered] to do *a*, it does not follow that the prima facie obligation has no moral effect on *M* at all" (Pietroski 1993, 509), and that "prima facie obligations may also *manifest* [my emphasis] themselves as residual obligations" (510). But while he makes use of dispositional notions like 'force' and 'manifestation' in his account of moral conflict, and while there are some limited affinities between our two accounts, he embeds his "diagnosis of moral conflict" in an account of moral principles on which they are *ceteris-paribus* generalizations (Pietroski 1993, 400, 499; cf. Brink 1994b, 216-17, 245). I reject this account of moral principles (see Chapter 5). He also suggests "that force talk is interchangeable with generalization talk" (Pietroski 1993, 499), which I deny, and he seems to deny that *prima facie* obligations are genuine obligations. However, we do agree on at least one much: that moral principles play the same explanatory role in normative moral theory that what he and (some) others take to be ceteris-paribus laws play in the special sciences (cf. Pietroski 1993, 499).

[17] Thus, Ewing's notion of "something positive in A" is closer to Ross's notion of a tendency than is "the notion that, other things being equal, A will be B, though this may be prevented by counteracting circumstances".

(and so from lighting when struck). Ross's view appeals to the fact that background conditions of the first sort, α-*conditions* as Stephen Mumford calls them (Mumford 1998, 86),[18] merely prevent the manifestations of dispositions, not their possession, in order to account for the fact that an outweighed *prima facie* obligation still makes a claim on an agent and that failing to fulfill that claim may give rise to residual obligations.

The idea is that when one *prima facie* obligation is *outweighed* by another in a particular case, this is merely an α-condition on the corresponding action's tendency to be obligatory overall. The corresponding action retains, unmanifested, its tendency to be obligatory overall. (Analogously, because my desk remains a massive body in earth normal gravity, it retains an unmanifested tendency to accelerate downward at a rate of 9.8 m/s^2, even though a number of factors, most notably the floor, prevent it from manifesting this tendency.) And on Ross's view, this is what accounts for the fact that outweighed *prima facie* obligations still make claims on agents, and thus for the fact that failing to fulfill those claims can give rise to residual obligations. (Somewhat analogously, we might say that the earth's gravity still makes a claim on my desk, even though it is resting on the floor.) Ross himself draws an analogy between his tendencies and mechanical tendencies, which are dispositions to move in particular directions with particular velocities. And the text will not bear a weaker reading. Ross's tendencies must be dispositions, and *prima facie* obligatoriness must be the tendency to be obligatory overall.

---

[18] The distinction between these two types of conditions is also drawn by Woolhouse (1973, 560-61).

4. Conflicts of Obligation: A Dispositionalist Account

This dispositional understanding of *prima facie* obligatoriness helps us to understand

the nature of conflicts of obligation. Moreover, Ross offers us additional resources with

which to construct an intuitive and plausible account of conflicts between genuine

obligations. Intuitively, what one ought (morally) to do rather than anything else just is

the most right thing one can do under the circumstances—be that to keep a promise, to

save a life, or what have you. Sometimes that is a particular action, or to perform a

particular action. But sometimes, as in the case of the surgeon, it is to perform one or the

other of two or more actions each of which is no less right than any other one could

perform. In other words, sometimes the most right thing one can do is to perform one

or the other of two or more maximally right actions. (In some cases, it may be to do

nothing, or even a lesser evil, but we can include such omissions and lesser evils within

these categories, as right actions.) But in either case, one's obligation overall is the most

right (or the least wrong) *course of action* one can pursue under the circumstances—no

matter how difficult it can be to figure out just what that is.

Although his terminology varies from mine, this is just Ross's view. He correctly

allows that moral rightness—which he takes to be moral fitness or suitability to a certain

situation—is a property that actions can have in different respects and in different

degrees (Ross 1939, 52-55, 60, 84, 100).[19] And he correctly distinguishes, first, between

_____

[19] In *The Right and the Good*, Ross takes the position that moral rightness is indefinable (Ross 1930, 12). But in *Foundations of Ethics*, he takes the position that moral rightness is partially definable, as moral fitness or suitability to a certain situation (Ross 1939, 53). On this view, for an action to be right in one respect is for it to be morally fitting or suitable to some aspect of an agent's circumstances (e.g., by fulfilling a promise the agent has made or by relieving another's

being obligatory and being maximally right and, second, between being maximally right and being obligatory overall. To be obligatory is to be right in (at least) one respect and to some degree, to be maximally right is to be no less right than any other action open to an agent in a situation, and to be obligatory overall is to be more right than any other such action. (In Ross's terminology, this is the distinction between being "right in one respect", being "right", and being "obligatory".)

Talk of degrees of rightness might sound odd to some, for there is a tendency among moral philosophers to speak as if the rightness of an action is an all-or-nothing affair. Similarly, because moral philosophers tend to use 'duty' and 'obligation' "in a sense in which 'my duty' or 'what I am obligated to do' is equivalent to 'the right thing to do', 'the best thing to do', 'what a virtuous man would do'" (Whiteley 1952-3, 54), talk of conflicts of obligation can sound absurd. Moreover, in his introduction to the new Oxford edition of Ross's *The Right and the Good*, Philip Stratton-Lake claims that "it makes no sense to say that a right act is one that is *most* right" and implicitly denies (a) that rightness admits of degrees and (b) that Ross believes that it does (Stratton-Lake 2002, xxxv). But neither Ross nor I say that *a* right act is one that is most right. Ross holds that *a* right action is one that is *no less* right than any other open to an agent (Ross 1930, 3-4; Ross 1939, 43-45, 52-53). And what we both hold is that *the* right course of action for an agent in any situation is the most right course of action open to her therein. All this makes perfect sense—at least once we distinguish (a) the property of being *a* right action

---

distress). And for an action to be more right than any other is for it to be more morally fitting or suitable to an agent's situation as a whole than is any other action open to her therein. But one needn't accept Ross's account of moral rightness as moral fitness or suitability to accept that moral rightness is one that actions can have in different respects and in different degrees.

in a situation, (b) the property of being *the* right action in a situation, and (c) the property *moral rightness.* 'Rightness' can refer to any one of these, but (c) and only (c) admits of degrees.[20]

Indeed, on Stratton-Lake's own account of what *moral rightness* is, it admits of degrees. He identifies rightness in the sense of (b) with "the property an act may have of having features that give us overall reason to do it" (xxii). By implication, he identifies *moral rightness* with *reasonableness*, which surely admits of degrees. At least some teleologists also allow that moral rightness admits of degrees, at least implicitly. David Brink argues that "a [moral] theory is teleological just in case it holds rightness to consist in maximal goodness " (1989, 215). By "rightness" he means "all-things-considered permissibility or obligation" (215), and by "maximal goodness" he means "the most value possible in the circumstances" (214). By implication, he claims that teleological moral theories just are ones that identify moral rightness with *goodness* and hold that *the* right course of action for an agent—an agent's obligation overall—in any situation is the most right (or best) course of action open to her therein.[21]

---

[20] Ross offers the following argument against anyone who might think him "dangerously lax in allowing degrees of bindingness": "let him suppose himself (*a*) to have promised to visit a sick friend whom he knows to be longing for his company, or alternatively (*b*) to have promised to go to the theatre with a party which he thinks will not miss him much if he does not go. In the absence of any other responsibility competing with either of these, each of them will be binding; but the second responsibility will be much more readily overridden by any competing responsibility, such as that of turning aside to help the victims of an accident" (Ross 1939, 100). (Ross uses 'bindingness', 'obligatoriness', 'rightness', and 'moral rightness' interchangeably, just as he uses 'responsibility' and '*prima facie* obligation' interchangeably.)

[21] It is a virtue of Stratton-Lake's and Brink's views that they preserve the intuitive idea that moral rightness admits of degrees. It is a further virtue that they purport to explain why it does. On the former view, moral rightness admits of degrees because some actions are more reasonable than others. On the latter, moral rightness admits of degrees because some actions are better

Now one might notice that if an action is right in one respect, then it will also be more right than any other one could perform if (but only if) nothing prevents it from being so. In other words, if an action is right in one respect, then it is also such as to be obligatory overall if nothing prevents it from being so. And even if such an action fails to be obligatory overall because an alternative action is more right than it is, it will remain right in one respect—i.e., obligatory. Moreover, if an action is right in one respect, then it also tends to be obligatory overall. Indeed, the tendency of an action to be obligatory overall just is its moral rightness under a different, overtly dispositional description.

Under one description (*viz.*, obligatory), an obligatory or *prima facie* obligatory action is one that is right in one respect, while under another (*viz.*, *prima facie* obligatory), it is one that tends (in a dispositional sense) to be obligatory overall. Thus, we can characterize a conflict of obligations in either of two ways. A conflict of obligations is a case in which a single action is right in one respect but wrong in another, or a case in which each of two or more mutually incompatible actions is right in one respect. Alternately, a conflict of

---

than others. If either (or both) of these views can be sustained, it would be an advance upon the minimal, abstract account that moral generalism offers of the metaphysics of morality, which is only a starting place, a foundation for further theoretical work. However, it is not clear whether these views can explain (rather than explain away) the phenomena of multiple obligations and conflicts between them. They would seem to suggest, respectively, that we have a *prima facie* obligation to perform any action there is some reason to perform (whether or not it is a moral reason) and that we have a *prima facie* obligation to perform any action that would be good in some respect. But, in each case, the view that we have *obligations* to perform such actions is counterintuitive. Thus, they seem to deny that *prima facie* obligations are obligations. One thing a teleologist might say here is that we have an obligation to perform F-actions just in case and because omissions of F-actions are wrong, and that omissions of F-actions are wrong just in case and because "some censure, blame, or punishment (as opposed, say, to praise or indifference) would be the best attitude to take toward" them (Brink 1989; see also Mill 1871, 5.14). A teleologist might also say that we have deontic reason to perform F actions just in case and because we have an obligation to perform F actions. Doing so might also allow the teleologist to explain why not all moral reasons obligate, because she could allow that actions can be right (possess moral rightness), and hence morally reasonable, without being obligatory (cf. page 61).

obligations is a case in which a single action has both a tendency to be obligatory overall and a tendency to be wrong overall, or a case in which each of two or more mutually incompatible actions has a tendency to be obligatory overall. Kant's case of the murderer at the door might be thought to illustrate the former sort of case, one in which a particular action is wrong *qua* lie but right *qua* saving of an innocent's life (8: 425), while the case of Abby and Andy illustrates the latter sort, one in which each of two (or more) mutually incompatible actions—my meeting Andy and my helping Abby—has a tendency to be obligatory overall (as does Sartre's case of the pupil torn between leaving to join the Free French and remaining to help his mother).[22]

The dispositional account of *prima facie* obligatoriness also allows us to characterize general obligations, such as the obligation to keep one's promises, and conflicts between them. Under one description (*viz.*, obligations), these are, or correspond to, action kinds that are right in one respect, while under another description (*viz.*, *prima facie* obligations) they are, or correspond to, action kinds that tend to be obligatory overall. Thus, the obligation to keep one's promises is, or corresponds to, the rightness, or *prima facie* obligatoriness, of promise-keeping *qua* (action) kind. Thus, we can characterize conflicts between general obligations in either of two ways. A conflict between general obligations is a case in which a single action is an instance of both an action kind that is right and an action kind that is wrong, or a case in which each of two or more mutually incompatible actions is an instance of an action kind that is right. Alternately, a conflict between

---

[22] One might wonder whether lying to Kant's murderer would, in fact, be wrong *qua* lie. That it would not be is something a moral holist can maintain without denying that other lies are wrong *qua* lies.

general obligations is a case in which a single action is an instance of both a *prima facie* obligatory kind (an action kind that tends to be obligatory overall) and a *prima facie* wrong kind (an action kind that tends to be wrong overall), or a case in which each of two or more mutually incompatible actions is an instance of a *prima facie* obligatory kind.

This account also allows us to characterize both what it is for one obligation to outweigh another and for two obligations to be equally weighty. For my obligation to help an injured person to outweigh my obligation to meet a friend who I promised to meet would be for my helping the former to be more right, or to have a greater tendency to be obligatory overall, than my meeting the latter. Similarly, for the general obligation to save an innocent life to outweigh the general obligation not to lie in Kant's case would be for lying to the murder to be more right *qua* saving of an innocent's life than it is wrong *qua* lie, or for it to have a greater tendency to be obligatory overall *qua* the former than it has to be wrong overall *qua* the latter. And for two or more obligations—or two or more instances of the same obligation—to be maximally weighty in a case would for each of them (or their instances to) be maximally right in that case, or to have tendencies to be obligatory overall of magnitudes equal to one another and greater than that of any other available alternative in that case.

Finally, we can characterize residual obligations dispositionally, as alternative manifestations of rightness *qua* (multi-track) disposition. Thus, when my obligation to help Betty, who is badly injured and needs my immediate assistance, outweighs my obligation to meet Andy, who I promised to meet, the rightness of my meeting Andy manifests, not in the obligatoriness overall of my doing so, but rather in the rightness of

my explaining to him why I was unable to keep my promise to meet him—that is, in my residual obligation to do so.

This dispositionalist account of conflicts of obligation is both intuitive and plausible. It also provides ready answers to the three questions that arise from the fact that I ought to help Abby rather than meet Andy, and that my meeting Andy would be both right and wrong. The questions are:

1. How can my meeting Andy have contrary moral properties? How can it be both right and wrong?

2. In virtue of what is my meeting Andy wrong?

3. In virtue of what is my helping Abby what I ought to do rather than anything else—i.e., my obligation overall?

The answer to (1) is that an action can have different properties in different respects, even contrary properties. My meeting Andy, for example, can be right in one respect and wrong in another; it can be right *qua* the keeping of a promise but wrong *qua* something else. The answer to (3) is that my helping Abby is obligatory overall in virtue of being more right than anything else I could do (including meeting Andy), or the most right course of action open to an agent in a situation. The answer to (2) is less straightforward, but not by much. The question here is really, in virtue of what is my meeting Andy wrong *overall*? It is not wrong overall in virtue of being other than the most right course of action open to an agent in a situation; and not every such course of action is wrong overall, as the case of the surgeon shows. Rather, my meeting Andy is wrong overall in virtue of being less than maximally right—i.e., because there is something else I could do that is more right (*viz.*, help Abby).

Finally, this account also gives an account of what an *absolute obligation* is, if there are any. Kant seems to have thought that the obligation to tell the truth is absolute (Kant 1797b). (He certainly denies that one may lie to a would-be murder at the door, even to save her intended victim's life.) Suppose that the obligation to tell the truth were absolute. What would be different about telling the truth if it were absolutely obligatory, rather than simply obligatory? It would be that instances of truth telling could not fail to be obligatory overall, or that nothing could prevent their tendencies to be so from manifesting (save, perhaps, their being physically impossible). Thus, what would make my telling the murderer the truth absolutely obligatory rather than simply obligatory would be that, in addition to its being right, or having a tendency to be obligatory overall, nothing could prevent it from being the most right course of action open to me, or from manifesting its tendency to be obligatory overall (save, perhaps, its physical impossibility). This is compatible with my lying to the murderer also being obligatory, but it does require that my doing so not be capable of being more right than is my telling him the truth. Given the difficulties inherent in establishing that any obligation satisfies these conditions, we ought to be skeptical about the existence of absolute obligations. But if there are any, they are, or correspond to, action kinds whose instances cannot fail to be obligatory overall, or to manifest their tendencies to be so.

5. Prima Facie Obligations and "Moral Obligations"

The dispositional account of *prima facie* obligations and *prima facie* obligatoriness (i.e., of what Ross calls "*prima facie* obligations" and "*prima facie* obligatoriness") that I advance

here differs from other, more familiar accounts.[23] For one thing, the "the notion of *prima facie* duties" is not, *pace* Ewing, "essentially the same thing" as "the conception of 'ought-making features'" (Ewing 1959, 111). Nor, also *pace* Ewing, does 'It is a *prima facie* duty not to tell lies' mean 'that X would be a lie is always a valid moral reason, though not necessarily a conclusive reason, against saying X' (Ewing 1959, 109-10). (Indeed, I will argue in Chapter 5 that the former does not even imply the latter.) For another, as obligations under a different description, *prima facie* obligations are neither right-making features, characteristics, or properties of actions (cf. McNaughton 1996, 79) nor moral reasons (cf. Urmson 1975, 113; McNaughton and Rawling 2000, 266).[24]

None of this is to deny that whenever an agent has a *prima facie* obligation to φ there is some factor that makes it *prima facie* right for her to φ; nor is it to deny that whenever an agent has a *prima facie* obligation to φ there is some moral reason for her to φ. But it is one thing to affirm that whenever an agent has a *prima facie* obligation to φ there is some factor that makes it *prima facie* right for her to φ and some moral reason for her to do so. It is quite another to identify *prima facie* obligations, moral reasons, and right-making factors, and yet a third thing to equate these three concepts. Indeed, it seems more natural to understand certain moral reasons—*viz.*, deontic, or obligating, ones—and right-making factors (be they ought-making features, right-making characteristics, or what have

---

[23] Stratton-Lake does indicate, though only in passing, that Ross's tendencies are dispositions, but he does not follow up on the idea—preferring instead to advance the view that "for an act to be *prima facie* right is not for us to have some kind of duty to do this act, but is for this action to have some feature that gives us a moral reason to do it" (Stratton-Lake 2002, xxxii-xxxiv).

[24] Nor are they moral principles (cf. Urmson 1975, 113).

you) as *grounding prima facie* obligations—i.e., as what make *prima facie* obligatory actions

and action kinds *prima facie* obligatory.

Moreover, the dispositional account of *prima facie* obligations avoids criticisms that

other accounts invite, especially those that understand *prima facie* obligations as, or in

terms of, moral reasons. In particular, it avoids the numerous criticisms that Bernard

Williams advances against what he calls "moral obligations". Williams recognizes "an

everyday notion of obligation," which he finds "ethically useful," but he objects to the

concept of "moral obligation", which belongs to what he calls "morality", or "the

morality system" (1985, 174). Although a discussion of Williams's critique of "the

morality system" is out of place here (as is a discussion of his arguments against what he

terms "ethical theory"), a discussion of what he finds objectionable about "moral

obligations" is apropos.[25]

One of the features that supposedly distinguishes "moral obligations" from

obligations is that, unlike obligations, "moral obligations" cannot conflict (Williams, 176).

But the everyday notion of an obligation that Williams finds ethically useful is itself a

---

[25] For responses to Williams's arguments against "ethical theory", see Nagel 1986; Nussbaum 2000. It is not always clear what Williams's positive views are, but he claims to accept what he calls "methodological intuitionism", the view that there is "a plurality of [ethical] first principles that may conflict" but no "explicit method or priority rules for resolving such conflicts," which is just Ross's view (1988, 183). (For Ross's slightly different understanding of "Intuitionism", see Ross 1939, 82; see also Ross 1930, 57-58; Ross 1939, 189.) Moreover, despite being critical of both Ross and cognitivist accounts of moral judgment like his (Williams 1965, 171, 175-76; Williams 1985, chap. 10), Williams himself takes a more or less Rossian line on conflicts of obligation (1981, 73-75), though he evidently does not recognize it as such. Unlike Ross, however, Williams is not a moral realist, save in the sense that he allows that some moral judgments can be true (cf. Williams 1988, 187-90; Williams 1985, 200). For some, of course, that is all it takes to be a moral realist.

moral notion: the notion of a moral obligation.[26] Thus, I think it best to understand

Williams's objections to "moral obligations" as objections to views about moral

obligations and the concept of a moral obligation, including those (like Kant's) that deny

the possibility of conflicts of obligation.

I agree with Williams that conflicts of obligation are possible, and the account I

offer of *prima facie* obligations holds them to be obligations and to be capable of

conflicting. Moreover, I offer an account of the metaphysics of such conflicts—one that

does not require resorting to non-cognitivism, or to any form of anti-realism, *pace* what

Williams has at times suggested (Williams 1965, 175; Williams 1966, 205). In contrast,

accounts that identify *prima facie* obligations with moral reasons or right-making factors

seem to suggest that they are not really obligations, and thus that conflicts of obligation

are impossible. It may be that some who offer such accounts do not mean to deny that

prima facie obligations are obligations, or that conflicts of obligation are possible (e.g.,

Brink 1994b, 215-16, 218-19). They may think that *prima facie* obligations are obligations

and that they are moral reasons. But since there are those who deny these things, those

who do not ought to welcome the dispositional account of *prima facie* obligations and the

account of conflicts of obligation that I advance here.[27]

---

[26] I see no reason to hold that all obligations are moral obligations, but the obligations Williams has in mind are moral ones, such as the obligation to keep one's promises. Unlike Ross, Williams distinguishes (moral) obligations from duties, which he suggests are connected to roles, positions, or relationships (Williams 1985, 7). I make no such distinction here.

[27] Stratton-Lake seems to deny that *prima facie* obligations are obligations (2002, xxxiii-xxxiv), as does Ross (Ross 1930, 18 & n. 1, 20; Ross 1939, 84-85), but Prichard evidently came to believe just the opposite (Prichard 2002, 287). For examples of those who explicitly deny the possibility of conflicts of obligation, see note 1.

Moreover, there are problems with identifying obligations with moral reasons. For Williams also raises a number of objections to "moral obligations" that can be summed up by saying that "the morality system" assumes that one is obligated to do is necessarily what one has the most moral reason to do (1985, 178-80, 185-86, 188-89). Once we allow that conflicts of obligation are possible, we deny that what one is *obligated* to do is necessarily what one has the most moral reason to do. But the question of whether what one is obligated *overall* to do is necessarily what one has the most moral reason to do is a different question. Williams's criticisms raise this undeniably substantive question, which accounts of *prima facie* obligations as, or in terms of moral reasons, tend to suggest is closed.

The question itself has two sides to it. First, are there moral reasons that do not obligate? That is, are there moral reasons to perform certain actions that do not make those actions either *prima facie* obligatory or obligatory overall? Can supererogatory and other, less extraordinary morally optional actions be favored by moral reasons?[28] Second, are there moral reasons that are capable of competing with obligations? And, if so, is it possible to have more moral reason to do something that one is not obligated to do than one has to do something that one is obligated overall to do. For example, might the balance of moral reasons favor your pursuing a unique opportunity to significantly further an important cause at the cost of breaking a promise to visit a friend, even if you have no

---

[28] There is also the question of whether and where we should draw the line between *moral* reasons and *ethical* reasons. Suppose supererogatory actions are reasonable. Are they supported by non-obligating moral reasons or by non-moral ethical reasons?

obligation to do the former?[29] If so, then the most right thing to do would not be the

morally best thing to do, and perhaps not what one ought to do in a moral sense of

'ought' that does not imply obligation.[30] There is also, of course, the further question of

whether what one has the most *moral* reason to do is necessarily that which one has most

reason to do (cf. Williams 1985, 187).

Here again, the real force of Williams's objection is to show that these questions are

open ones. The dispositional account of *prima facie* obligations leaves open these and

other questions about what relations obtain between moral and other reasons, on the one

hand, and obligations (*prima facie* and overall), on the other. Thus, unlike accounts that

understand *prima facie* obligations as, or in terms of, moral reasons, it does not invite

Williams's objection.[31]

---

[29] The case is Williams's, but the question is more complicated than he suggests. He allows
that if your friend favors the cause and your support of it, and if she would release you from the
promise if you could contact her, then you might not have an obligation to keep the promise
(Williams 1985, 180). For my part, I do not know what to say about such a case. My intuition is
that you could be morally justified in breaking the promise, and that it could be favored by the
balance of moral reasons, but it is not clear to me that you would have no obligation to keep it.
Certainly, you would owe your friend an apology if you broke your promise. Nor is it clear to me
what to say about whether your friend might have an obligation to release you from the promise
if the cause is just. Suppose that she does. Might you have no obligation to keep the promise?
Again, it is not clear to me what to say here, but my intuition is that one has an obligation to keep
the promise, even if the morally reasonable thing to do is to break it.

[30] Or, if we should limit the moral to the obligatory, the *right* thing to do would not be the
ethical thing to do.

[31] It also provides common ground—in the form of shared (or at least sharable)
concepts—for those who disagree about what the answers to such questions are. This might
make such disagreements more tractable, if for no other reason than that it relocates disputes
over the existence of moral obligations and *prima facie* obligations, and over the viability of these
concepts, to disputes over what relations obtain between moral obligations (*prima facie* and overall)
and reasons (moral and otherwise).

There is one last feature of Williams's "moral obligations" that *prima facie* obligations might be thought to share, even on the dispositional account. Williams objects to the view that "particular obligations" must be backed by "general obligations". It is not entirely clear to me which of three distinctions he has in mind, that between the general and the specific (e.g., between lies and lies told to mothers about their cooking), that between a kind and its instances (e.g., between lies *qua* kind and individual lies), or that between the universal and the singular (e.g., between all lies and a lie) (1985, 175, 180-82, 187 esp. note 13). The cases he has in mind, at any rate, are ones like those in which one is *confronted* by someone in danger, or at risk. And he clearly wants to deny that if one is obligated to try to help such a person, one is necessarily "under an obligation to assist all people at risk, or to go around looking for people at risk to assist" (181).

Now I do not share Williams's confidence that we are not under an obligation to assist everyone at risk, but neither the dispositional account of *prima facie* obligations nor moral generalism requires that obligations be general rather than specific—although I do suspect that most, if not all, fundamental (or basic, or underived) obligations are quite general. Thus, they both allow that we may only be under an obligation to assist those at risk whom we could assist with minimal risk to ourselves (cf. 181). But both the dispositional account of *prima facie* obligations and moral generalism might be thought to require both that every particular *prima facie* obligatory action be an instance of a general *prima facie* obligation, or of an action kind that is *prima facie* obligatory, and that every instance of such an obligation or kind be *prima facie* obligatory.

Ross certainly suggests that they require the former when he claims that *prima facie* obligatory actions are *prima facie* obligatory in virtue of being instances of morally

significant kinds. And he seems to suggest that they require the latter when he writes that the only way to save the authority of commonsense moral rules such as 'tell the truth', 'injure no man', and so on, "is to recognize them not as rules guaranteeing the rightness of any act that falls under them, but as rules guaranteeing that any act which falls under them tends, so far as that aspect of its nature goes, to be right, and can be rendered wrong only if in virtue of another aspect of its nature it comes under another rule by reason of which it tends more decidedly to be wrong" (Ross 1939, 313). Insofar as Ross sometimes uses 'right' to mean 'obligatory overall', this might be taken to suggest that every instance of a *prima facie* obligatory action kind must itself be *prima facie* obligatory.[32] Moreover, if the dispositional account of *prima facie* obligations and moral generalism do require that every instance of a *prima facie* obligatory kind be *prima facie* obligatory, then it would seem that they are incompatible with moral holism. For it is hard to see how, say, promising to φ could fail to make φ-ing *prima facie* obligatory if every instance of promise-keeping must be *prima facie* obligatory.

In Chapter 5, I will argue that neither moral generalism nor the dispositional account of *prima facie* obligations requires either that every particular *prima facie* obligatory action be an instance of a *prima facie* obligatory action kind or that every instance of such a kind be *prima facie* obligatory. Moreover, I will argue that moral dispositionalism cannot

---

[32] Similarly, Pietroski suggests that we think of moral principles as either *ceteris paribus* generalizations (or laws) of overall obligatoriness or universal (or strict) generalizations (or laws) of *prima facie* obligatoriness (1993, 501). He also suggests that we understand talk of actions being instances of moral kinds as talk of their falling under moral generalizations, or laws (492). I reject these accounts of moral principles and of what it is to be an instance of a moral kind for the same reason I reject the view that moral principles are universal moral statements or truths (Chapter 1). Indeed, I reject all accounts of moral principles as, or in terms of, deontic laws (Chapter 5).

only reconcile moral holism and moral generalism, but also explain them both. But first,

I consider the substance and merits of moral holism.

CHAPTER 3

MORAL HOLISM

Moral holism maintains that what is a moral reason to φ in one case may not be one

in another, and may even be a moral reason not to φ given the right context. It holds that

the moral polarity or valence, as it were, of a moral reason is not fixed independently

of—and so may be altered by—factors other than itself. And some "moral particularists"

maintain that moral holism motivates skepticism about the existence of and need for

moral principles, along with skepticism about the viability of principle-based approaches

to ethics and moral theory (see page 130). But before we can assess either the merits or

implications of moral holism, we need a satisfactory understanding of both it and its

contrary, *moral atomism*. In this chapter, I explicate moral holism and moral atomism, and

distinguish the former from *particularism in the theory of moral reasons*, which is a special case

of *holism in the theory of reasons*. I also relate moral holism to views held by two advocates

of systematic, principled moral theory, Shelly Kagan and Thomas Hurka, and show why

they, too, are moral holists.

1. <u>Moral Holism</u>

Moral holism is often (and most intuitively) stated as a view about moral reasons.[1] So stated, it claims that what is a moral reason to φ in one case may be no moral reason to φ or (even) a moral reason not to φ in another, but it understands moral reasons in a particular way. It understands them as contributory or *pro tanto* reasons, rather than as decisive or necessarily overriding ones. And it understands them as factors that contribute to the rightness and wrongness of particular actions, but whose individual contributions are not necessarily determinative of whether an action is right or wrong *overall*. Thus, to say that there is moral reason to φ is to say that there is a (or some) moral reason to φ, or that φ-ing is *pro tanto* right. It is not to say that there is sufficient moral reason to φ, or that φ-ing is *the* morally reasonable or *the* morally right thing to do under the circumstances. There could be more moral reason not to φ than there is to φ, or there could be more moral reason to ψ than there is to φ. Contributory or *pro tanto* reasons can be outweighed in particular cases without losing any of their normative force therein—and *a fortiori* without being voided or shown to be of merely epistemic relevance therein.[2]

---

[1] Forms of moral holism are advanced or endorsed in Dancy 1993, 60-62; Dancy 2000, 130-31; Dancy 2004, 73-80; Kagan 1998, 307; Kagan 1988, esp. 16-18 & n. 7, 20-21 & n. 8; McNaughton 1988, 192-93; McNaughton 2000, 285; McNaughton and Rawling 2000, 257-58; Little 2000, 281; Ross 1939, 109-12; Shope 1965; Sinnott-Armstrong 1999, 3; Williams 1985, chap. 10; Williams 1995, 57.

[2] Most moral holists will, I expect, hold that whether φ-ing is *the* morally reasonable or right thing to do depends on whether there is, on balance, more moral reason to φ than there is to do anything else—allowing that more than one action may be equally morally reasonable or right for an agent at the same time. But they need not identify what one has most moral reason to do with what one is morally obligated to do (cf. Williams 1985, 178-79). Indeed, a moral holist may reject

This conception of a moral reason is a familiar one, with roots in Ross's account of *prima facie* obligations. It understands (deontic) moral reasons (i.e., reasons that obligate) as contributory *right-making factors*.[3] (For convenience, I use the term 'right-making factor' to cover both right-making and wrong-making factors.) Ross holds that every *prima facie* obligation "rests on" a morally significant circumstance, or fact; for example, that one has promised to ϕ, or that there are other beings in the world whom one can help (1930, 20-21). And these morally significant facts are naturally understood as contributory right-making factors, as is suggested by Ross's explicit concern with the question of "what makes right acts right," his pluralism about the number of basic *prima facie* obligations, and his belief that there can be (genuine) conflicts between such obligations. They are also naturally understood as (contributory or *pro tanto*) moral reasons, as they are the kinds of facts one would naturally take to be moral reasons, or reasons from the moral point of view (cf. Frankena 1973, 113-14).

---

deontic notions like obligation all together. (Such a holist might prefer to speak of "ethical" rather than "moral" holism, and may eschew terms like 'right' and 'wrong' for terms like 'good' and 'bad', or 'virtuous' and 'vicious', when speaking of the ethical properties or qualities of actions.) Moral holism is also neutral on whether what one has most moral reason to do is always morally obligatory (i.e., on whether there are non-obligating moral reasons) and on whether what one has most moral reason to do is always what one has most reason to do (i.e., on whether the moral thing to do is always the reasonable thing to do).

[3] Although it has roots in Ross's work, this conception of a moral reason is in many respects an ecumenical one. First, it can be embraced by those who do not share Ross's pluralism about the number of fundamental moral obligations. Indeed, consequentialists and deontologists alike sometimes represent their dispute as one over whether there is just one fundamental *prima facie* obligation or *pro tanto* moral reason—namely, to promote the good—or more than one (e.g., Frankena 1973, 43-48; Kagan 1989, 16-18). Second, this general conception of moral reasons can be embraced by deontologists who do not share Ross's view that there are no absolute moral obligations. For, as a general account of moral reasons, it is compatible with the existence of absolute obligations. One may even accept it if one believes that all moral reasons are, as a matter of fact, determinative of whether an action is right or wrong overall. This ecumenical character is one of the major strengths of this conception of a moral reason, as is its institutive appeal.

Although many who share this conception of a moral reason allow that the strength, or weight, of such a reason may vary from one case to another, it would seem that many regard such reasons as factors that always make the same kind of contribution (positive or negative) to the rightness (or wrongness) of actions.[4] Moral holists, however, claim that even the moral polarity or valence, as it were, of a moral reason may vary from one case to another. That is, they claim that the same (type) of factor may be a moral reason to φ (a right-making factor) in one case, no moral reason (a morally indifferent factor) in a second case, and even a moral reason not to φ (a wrong-making factor) in a third.

To explain how and why a moral reason's valence may vary from one case to another, moral holists appeal to the effects that other factors may have on that valence. Specifically, they claim that the valence of a moral reason in a particular case may be affected by—and is thus not fixed independently of—other factors. Now, many would allow that the *weight* of a moral reason in a particular case may be affected by other factors. That is, many hold the following *weak dependent-contributions thesis*:

> The moral *weight* of a factor *f* in circumstance *c* may be affected by factors other than *f* itself.[5]

Anyone who holds this thesis might be called a moral holist. But the moral holists I have in mind also accept the following *strong dependent-contributions thesis*:

---

[4] See, e.g., Ewing 1959, 110-11, 126; Frankena 1973, 26; Urmson 1975, 112-13; Brink 1994b, 219; Audi 2002, 52. Brandt does not speak in terms of moral reasons, but he does assume that "obligation-creating factors" always create *prima facie* obligations (Brandt 1959, 368).

[5] Those who accept this thesis *eo ipso* reject what Kagan calls "the assumption of independent contributions," which is the assumption "that the size of a [morally relevant] factor's contribution is *independent* of everything other than the value of the factor itself" (1988, 17). I say more about this below.

> The moral weight *and valence* of a factor *f* in circumstance *c* may be affected
> by factors other than *f* itself.

This thesis purports to explain why a factor that is a moral reason to ϕ in one case may be

no moral reason to ϕ or a moral reason not to ϕ in another: the moral valence of a factor

may vary from one case to another because other factors may affect its moral valence in

particular cases.

There is a degree to which moral holism is (I think) uncontroversial. Few, if any,

would say that its being Tuesday is morally relevant in its own right. But if I have

promised to ϕ on Tuesday, its being Tuesday may be of considerable moral relevance.

While it seems plausible that whatever moral relevance its being Tuesday has in a

particular case derives solely from the relations in which it stands to other morally

relevant factors (e.g., the fact that I promised to ϕ on Tuesday), the moral holism I have

in mind is not limited to factors of derivative moral relevance, or derivative moral

reasons. It is a view about factors that make non-instrumental and otherwise non-

derivative contributions to the rightness (and wrongness) of particular acts—factors

whose moral relevance does not derive solely from the relations in which they stand to

other morally relevant factors. But neither does it deny that factors have normal, default,

or natural moral valences (cf. Dancy 1999b, 27; Dancy 2004, 112-13).[6] The strong

dependent-contributions thesis does not say that the moral valence of *f* in *c* is wholly or

even mostly dependent on factors other than *f* itself. It claims only that the moral valence

of *f* may be affected by factors other than *f* itself. Moreover, moral holists maintain that if

---

[6] One can, of course, imagine a more radical moral holism that denies this.

a factor *f* has a moral valence in one case that is different from the moral valence it has in another, then there must me an explanation of this—an explanation in terms of the effects that one or more other factors have on its moral valence.

Implicit in both the strong dependent-contributions thesis and the weak dependent-contributions thesis is a tripartite distinction between *moral reasons*, *active background factors*, and *inert background factors*.[7] Moral reasons on this view are factors that contribute (positively or negatively) to the rightness of an action; they are, that is, *contributory factors*. Active background factors are ones that may affect the contribution made by moral reasons in particular cases—i.e., factors that may affect the weight or valence of contributory factors. Active background factors *as such* do not *contribute* to the rightness of an action. But they do *affect* the rightness of particular actions, by affecting the contributions other factors make to the rightness (and wrongness) of those actions. They may enable other factors to have particular moral valences or prevent them from doing so. They may even reverse the moral valences of other factors (e.g., from positive to negative).[8] Inert background factors are ones that have no effect, direct or indirect, on the rightness of actions.

---

[7] This distinction was first introduced by Dancy, as one between different properties: (1) properties in virtue of which a right action is right, (2) properties the presence or absence of which affects or would affect the way in which properties of the first sort make an action right, and (3) properties that have no such effect (Dancy 1993, 55). But it can also be expressed in the idiom of facts, as the distinction between (1) facts in virtue of which a right action is right, (2) facts the presence or absence of which affects or would affect the way in which facts of the first sort make an action right, and (3) facts that have no such effect (56).

[8] See, e.g., Kagan 1998, 307; Dancy 1999b, 26, 32; Sinnott-Armstrong 1999, 2-8. The distinction between moral reasons or contributory factors and active-background factors is analogous to that between causes and enabling conditions.

The claim that active background considerations *as such* do not contribute to the rightness of an action might require some clarification. Neither the strong nor the weak dependent-contributions thesis requires that active background factors be reasons (moral or non-moral) in *c*. They allow that the weight or valence of, say, a moral reason to φ in *c* may be affected by factors that are not themselves either reasons to φ or reasons not to φ in *c*. (They need not be reasons [not] to ψ in *c* either.) Thus, suppose that I borrow a book from you and that you subsequently tell me to keep it. The fact that I borrowed the book from you ceases to be a reason for me to return it to you, because you subsequently gave it to me. But the fact that you subsequently gave it to me *need* be neither a reason for me to keep it nor a reason for me not to keep it. Indeed, we can easily imagine cases in which the fact that you subsequently gave the book to me is simply irrelevant to what, if anything, I should do with it. Thus, an active background factor need not itself be a reason. But we can also easily imagine cases in which the fact that you subsequently gave it to me is a reason for me not to return it to you or, more generally, to keep it.[9] In such cases, the fact that you subsequently gave it to me is both an active background factor and a reason. *Qua* reason, it contributes to the rightness of various (possible) actions. But *qua* active background factor, it does not. *Qua* active background factor, it affects but does not contribute to the rightness of various (possible) actions.

It may help to consider another example. Here is a schematic one of Jonathan Dancy's that serves to illustrate moral holism, the strong dependent-contributions thesis,

---

[9] We can also easily imagine cases in which the fact that you gave it to me is a reason for me to give it to someone else. What the holist will say, of course, is that whether a particular case is one rather than another of these three kinds of cases depends (at least in part) on other factors, which need not be reasons for me to do one thing or another with the book.

and the tripartite distinction between reasons, active background factors, and inert

background factors:

> The fact that an action will give pleasure can be a reason for doing it or
> for approving of it when done. But it can also be a reason for
> disapproving of it. If I tread on a worm by mistake, my action is perhaps
> morally indifferent. But if I tread on it with pleasure or to give you
> pleasure, my action is the worse for it. We should not say, as I once heard
> Professor Hare argue, that there is at least this to be said for the action,
> that it gave me (or you) pleasure, though of course we must condemn the
> action overall because the value derived from the pleasure is overwhelmed
> by other disvalue. (Dancy 1993, 56)

Dancy's view and the one he attributes to R.M. Hare treat the fact that ϕ-ing would give

someone pleasure and the fact that the pleasure would be taken in the harm done to the

worm differently. The latter regards (or at least appears to regard) the fact that ϕ-ing

would give someone pleasure as an *invariant moral reason*—a factor that must be a moral

reason to ϕ (a right-making factor) in any case, even one in which ϕ-ing consists in

harming or even killing another sentient being (Dancy 2000, 137).[10] Consequently, it

regards the further fact that the pleasure would be taken in the harm done to the worm as

morally indifferent, or as an inert background factor.

Like moral holists generally, Dancy allows that there may be invariant moral

reasons, but like many other moral holists, he does not believe the fact that ϕ-ing would

---

[10] Strictly speaking, the claim Dancy attributes to Hare is much weaker than this. It is
simply that the fact that ϕ-ing would give someone pleasure in the case described is a moral
reason to ϕ therein. This claim is clearly consistent with moral holism. One might also note that,
given his prescriptivism, it is odd to think of Hare as a moral atomist (the opposite of a moral
holist), because (like moral holism) moral atomism is a view about right-making factors. But
Dancy intends this example as an illustration of "how the logical behaviour of reasons is like that
of morally relevant properties" (1993, 56); that is, he intends it as an illustration of how holism in
the theory of reasons (see below) is like moral holism. I use it to illustrate the latter view.

give someone pleasure is such a reason.[11] Dancy's view regards the fact that φ-ing would

give someone pleasure as a *variant moral reason*—a factor that is a moral reason to φ in

some (morally possible) cases, but also no moral reason to φ or a moral reason not to φ in

others. Dancy takes it to be a moral reason not to φ here, a wrong-making factor. And

he takes this to be because the pleasure would be taken in the harm done to the worm.

That is, he thinks the fact that the pleasure would be taken in the harm done to the worm

is an active background factor that affects—in this case, reverses—the moral valence of

the fact that φ-ing would give someone pleasure—i.e., the effect it has on the rightness of

φ-ing. Moreover, he does not take the former to be a reason, at least not here; nor does

he take it to be part of a "complete reason" (see below); nor does he take the relevant

reason to be a more specific one, such as that φ-ing would give someone *sadistic* pleasure

(see below). Rather, he takes the relevant reason here to be (simply) the fact that φ-ing

would give someone pleasure. And he takes the further fact that the pleasure would be

taken in the harm done to the worm to be an active-background factor—a factor that

affects the moral valence (or weight) of another factor but that need not itself be a reason

(not) to φ (or a part thereof).

---

[11] As McNaughton writes, "It is important to distinguish three different positions here. First, there is Ross's view that if a characteristic matters in one place it *must* [my emphasis] matter everywhere. Second, there is the weaker view that there *must* [my emphasis] be at least some characteristics which always matter, even if there are others which matter only sometimes. The particularist rejects these positions. Third, there is the much weaker view that there *may* be some characteristics which, as a matter of fact, always count in the same way. The particularist need have no quarrel with this position" (McNaughton 1988, 201). It is not clear to me whether McNaughton's "particularist" is a moral holist, a particularist in the theory of moral reasons (see below), one who rejects moral principles and principled-based approaches to ethics and moral theory, or what.

One might wonder how a moral holist can allow that there may be invariant moral reasons. To be sure, moral holism is a view about right-making factors as such, so it is true of all of them if it is true of any of them. But its modality is as weak as can be. Moral holism is merely a view about what is *conceptually* possible. It is the view that it is conceptually possible for a right-making factor to lose or reverse its moral valence from one case to another as the effect of changes in background factors. That is, it is a view about what the concept of a morally relevant factor—or, better, the concept of something that is relevant, or makes a difference, to what one ought morally to do—entails. As Dancy puts it, it is a view about *the logic* of such factors, or that of the right-making relation (Dancy 2000, 137; Dancy 2004, 79).[12] And as such, it implies nothing about whether any particular factor's moral valence is variant or invariant. It only implies that if a particular factor's moral valence is invariant, this is not because right-making factors are, as a conceptual matter, valence-invariant.

This leaves open numerous possibilities. For one thing, a particular factor's moral valence may just be invariant. To adapt an analogy of Dancy's, any human *qua* human can run a four-minute mile; it is conceptually possible for any human to run a four-minute mile. But this is consistent with the fact that Smith cannot (Dancy 2000, 136). If Smith cannot run a four-minute mile, it isn't because it is conceptually impossible for a human to run a four-minute mile; it is because of some fact about Smith other than that

_____

[12] Similarly, holism in the theory of reasons and particularism in the theory of moral reasons (see below) are views about the logic of the favoring relation (Dancy 2004, 79) and of "the very logic of moral reasons" (Dancy 2004, 76), respectively. They are all on a par with Hare's "thesis of universalizability", which Hare calls a "logical thesis", by which he means "a thesis about the meaning of words, or dependent solely upon them" (Hare 1963, 30).

he is a human (e.g., that he has a bad knee). For another thing, the invariance of a particular factor's moral valence may be necessitated by the fact that it falls under some *other* concept. Margaret Little claims, "Thick moral features differ from nonmoral ones precisely because, so identified, they are guaranteed of carrying a given valence of moral significance (part of what it is to *count* as a moral feature, to earn the status as a moral feature so identified, is to count as a moral reason of a given direction)" (Little 2000, 289). One might well contest this, but it is consistent with moral holism. Any human *qua* human may be a bachelor; it is conceptually possible for any human to be a bachelor. But this is consistent with the fact that it is conceptually impossible for Sally to be a bachelor. If it is conceptually impossible for Sally to be a bachelor, it isn't because it is conceptually impossible for a human to be a bachelor; it is because of some fact about Sally other than that she is a human together with some fact about the concept of a bachelor (*viz.*, that she is a woman and that a bachelor is an unmarried man). Thus, it is quite possible for a moral holist to allow that there may be invariant moral reasons, and even that there are many.

Indeed, moral holism is sufficiently weak that a moral holist could deny that there are any *variant* moral reasons, just as one could hold that it is conceptually possible for a human to run a *one*-minute mile but that no human could run a mile that fast. To be sure, many moral holists believe that there are valence-variant right-making factors: that some variant moral reasons exist. But this view—call it *ontological holism*—is a stronger view than moral holism, and it is possible for a moral holist to hold that all *existent* right-making factors are valence-*in*variant. That is, a moral holist could be an *ontological atomist*. Indeed, a moral holist could—I suppose—be a *metaphysical atomist*. That is, a moral holist

could deny that variant moral reasons are *metaphysically* possible. For example, a moral

holist could be a hedonistic act utilitarian who *identifies* the property of rightness (i.e., of

being morally permissible or obligatory overall) with the property of being such as to

promote a maximum of pleasure. On this view, that φ-ing would promote pleasure is an

invariant moral reason; there is no metaphysically possible case, and hence no morally

possible case, in which it does not contribute to the rightness of φ-ing.[13] But if she

regards this identity as a synthetic or an *a posteriori* one, rather than analytic or *a priori*, then

her view is consistent with moral holism, for moral holism is merely a view about what is

conceptually possible. Moral holism leaves open the possibility that variant moral reasons

are metaphysically impossible, just as the view that nothing in the concept of water entails

that water is H_2O leaves open the possibility that water is necessarily H_2O.[14] (That is,

moral holism no more entails that variant moral reasons are metaphysically possible than

the fact that 'Water = H_2O' is not a *conceptual* truth entails that the identity of water and

---

[13] A weaker, non-reductive version of this view would hold that the property of rightness is constituted by, but not identical to, the property of being such as to promote a maximum of pleasure. On this weaker view, hedonistic act utilitarianism is true in the actual world and (trivially) in all morally possible worlds, but not in all possible worlds. It is true, but not necessarily true: rightness might have been constituted by some other property. This is an atomistic view, but only an ontologically atomistic one, for it entails only that there *are* only invariant moral reasons.

[14] Or, if one prefers, it leaves open the possibility that there could be no variant moral factors in the same way that the A-intension of the sentence 'All water is H_2O' being contingent and *a posteriori* leaves open the possibility that water = H_2O, and hence the possibility that the C-intension of the sentence 'All water is H_2O' is necessary *a posteriori* (Jackson 1998, 84-86). Jackson argues we should not distinguish *in kind* between conceptual necessity and metaphysical necessity, in part on the grounds that we don't need two kinds of necessity to explain the phenomena of the necessary *a posteriori* (1998, 67-74). The moral holist needn't take a position on this, because—at heart—her position is one about what we are entitled to assume prior to and for the purposes of inquiry. She could restate her view as one about what is epistemically possible or contingent *a priori*, and indeed it may be that some moral holists think of their view in such terms.

H_2O is metaphysically contingent.) Thus, not only could a moral holist deny that there *are* any variant moral reasons, she could deny that there *could be* any such reasons—just as one who denies that the identity of water and H_2O is a conceptual necessity could deny that water is or could have been anything but H_2O.

Although it might seem puzzling that a moral holist can deny that there are any variant moral reasons, or even that there could be none, this is no more puzzling than that one could hold both that nothing in the concept of a moral obligation entails that moral obligations are absolute obligations and that there is a single, absolute moral obligation (e.g., to maximize utility, or to respect humanity, or to love God and one's neighbor), or even that there is a single such obligation necessarily. One might well hold that there is nothing in the concept of an obligation that is inconsistent with a plurality of non-hierarchically or non-lexically ordered obligations and yet hold that there is just one (trivially absolute) obligation. A moral holist is merely one who holds valence-variant right-making factors are *conceptually* possible, not one who holds that they are possible *simpliciter*. And it is this surprisingly weak thesis that some moral particularists maintain motivates skepticism about (a) the existence of and need for moral principles and (b) the viability of principle-based approaches to ethics and moral theory.

2. Moral Atomism

If a moral holist is one who holds that valence-variant right-making factors are conceptually possible, then a *moral atomist* is one who denies this, or holds it to be a conceptual truth that a factor's moral valence cannot vary from one case to another. If holism is indeed uncontroversial as a thesis about factors whose moral relevance is wholly

derivative of the relations in which they stand to other morally relevant factors, then

moral atomism is the moral-theoretic analogue of the view that "all judgments of intrinsic

value are . . . universal" in the sense that "a judgment which asserts that a thing is good in

itself, if true of one instance of the thing in question, is necessarily true of them all"

(Moore 1903, 27; cf. Korsgaard 1996, 251, 257). It is, in other words, the view that, as a

conceptual matter, *intrinsic* moral relevance is *invariable* intrinsic moral relevance.

Brandt is evidently a moral atomist, for he claims that "if we are ever in a position

to say that we have a prima facie obligation to do *A* on account of *F*, we are in a position

to say that *always* there is a prima facie obligation to do *A* on account of *F*, for any

evidence we have against the universal statement counts against our right to speak of a

prima facie obligation being present at all" (Brandt 1959, 368 n. 13). And one might

think that Moore is a moral atomist simply on the strength of his *value* atomism. For it is

not clear why one would be an atomist about intrinsic value but not about intrinsic moral

relevance. But Moore also implies that moral atomism is true when he claims that "'right'

does and can mean nothing but 'cause of a good result', and is thus identical with

'useful'" (Moore 1903, 147). For if that were true, then not only would it make no sense

to say that something other than an action's usefulness might contribute to its rightness, it

would also make no sense to say that an action's usefulness might fail to contribute to its

rightness. In other words, this view implies that the holist's thesis is, literally, nonsensical.

Hare, too, implies that moral atomism is true when he claims that moral judgments

are universalizable in the sense that if one says that an action is right and thinks that one

is entitled to do so because it is F, then one is committed, as a matter of logic, to the view

that every F action is right (Hare 1963, 10-12).[15] If that were true, one could not

consistently claim that one action is right simply because one promised to perform it but

that a second is not even *pro tanto* right because, although one promised to perform it,

that fact makes no contribution to the rightness of doing so because the action would be

a nefarious one. In other words, Hare's view implies that moral holism is logically false.

Brandt, Moore, and Hare demonstrate that one can be a moral atomist without

saying either that variant moral reasons are conceptually impossible or that intrinsic moral

relevance is, as a conceptual matter, invariable intrinsic moral relevance. And moral

particularists evidently believe that moral atomism is not only quite widespread, but also

that their moral holism distinguishes them from most, if not all, moral philosophers, past

and present.[16] But it is difficult, if not impossible, to determine just how widespread

moral atomism really is. One reason for this is that one cannot tell whether someone is

an atomist from the mere fact that she advances a moral theory according to which there

are only invariant moral reasons. One could hold that all moral reasons are invariant ones

yet deny that this is conceptually (or even metaphysically) necessary. Thus, if a moral

theory recognizes only invariant moral reasons, one cannot assume that this is because its

proponents thought variant ones conceptually, or even metaphysically, impossible. They

could simply believe, on reflection, that there are only invariant moral reasons.

---

[15] As a claim about the meaning of the word 'right', which it is (Hare 1963, 30), this claim is simply false. Saying that John's returning Jane's gun to her was right and thinking that one is entitled to say it was right (simply) because Jane owned the gun does not commit one to the view that every returning of a gun to its owner is right. One could consistently do so while holding that some such returnings might have other features that make them wrong, despite their being such returnings.

[16] See, e.g., Dancy 1993, 66; McNaughton 1988, 191-94; Little 2000, 279-80.

Another reason for this—for why it is difficult to tell just how widespread moral

atomism is—is that some who appear to be moral atomists may well be moral holists.

Some of Dancy's critics favor including relevant contextual factors—i.e., active

background factors—*in* moral reasons, as parts thereof.[17] As Joseph Raz formulates this

view, it distinguishes between *complete reasons* and *parts of (complete) reasons* and counts as the

latter both factors that contribute to the rightness and wrongness of actions (e.g., that I

borrowed a book from you) and factors that affect the nature of the contribution made

by those contributory factors (e.g., that you had a right to posses the book or that you

stole it). Thus, on this view, a moral reason to ϕ is the totality of factors (including

contextual factors) relevant to whether an agent has a moral reason to ϕ in *c*.

This is an atomistic view of moral reasons, but it neither is nor entails moral

atomism. In fact, it entails moral holism. If a moral reason to ϕ is indeed the totality of

factors (including contextual factors) relevant to whether an agent has a moral reason to ϕ

in *c*, then it follows that a moral reason to ϕ in one case must be a moral reason to ϕ in

any case. Thus, this is an atomistic conception of moral reasons. But inasmuch as this

view holds that both contributory factors and active background factors are parts of

(complete) reasons, it entails moral holism. For moral atomism disallows active

---

[17] Both Raz and Hooker favor this as a view about reasons in general (Raz 2000, 58-61, 67-69; Hooker 2000, 6-11, 12-15). For simplicity, I consider this view as it applies to moral reasons only. However, my claims apply *mutatis mutandis* to their general view. Crisp may fall into this category, as well (see Crisp 2000, esp. 33-44 & n. 46). If not, his view must be one of the more radical ones I describe in note 18. Interestingly, Raz and Hooker have very different views about what the consequences of this view are. Hooker sees it as a way of defending the view "that morality is composed of an irreducible plurality of principles" (Hooker 2000, 2), while Raz claims to adopt this view is to "reject the thought that morality or any other domain of practical rationality consists in principles" (Raz 2000, 77).

background factors: according to atomism, all background/context is inert

background/irrelevant.[18]

What is going on here seems to be this: some want to reject both the (broadly

Rossian) conception of (deontic) moral reasons as contributory right-making factors and

holistic conceptions of moral reasons, but accept both contributory right-making factors

and holistic conceptions thereof.[19] Given this, we need to distinguish between moral

holism and *particularism in the theory of moral reasons*. The former is a view about right-

making factors as such, and a view about moral reasons only to the extent that moral

reasons are such factors. The latter, which is a special case of *holism in the theory of reasons*

(Dancy 2000, 130), is a view about moral reasons *qua* reasons.[20] But while failing to

---

[18] We might distinguish this view from two possible alternatives. The first identifies right-making factors with moral reasons, holds that these factors/reasons are valence-invariant complexes of simpler factors, and holds that these simpler factors are themselves invariantly morally indifferent. In other words, it holds that right-making factors/moral reasons are (quasi-Moorean) organic unities whose parts have no moral relevance of their own either inside or outside of such unities. This would be an atomistic view of right-making factors. The second possible alternative denies that there is a real distinction between right-making factors and active background factors. This would be an even more radical moral holism than the one under discussion. For it denies there is any real distinction between the ways in which different factors (that φ-ing would give someone pleasure, that the pleasure given would be taken in the harm done to a worm, that I borrowed a book from you, that you did not steal it, that I am able to return it to you, etc.) are relevant to the rightness of φ-ing (e.g., stepping on the worm or returning the book to you).

[19] One reason for distinguishing between right-making factors and moral reasons is that on some views the reasons one has to φ and the factors in virtue of which φ-ing is right can come apart, at least in some cases. For example, both Raz and Slote hold that an agent's character or motives can contribute to (or even determine) the rightness of her acts, but that they are not reasons for her to do them (Raz 2000, 77-78; Slote 2001, 38-50).

[20] Moral holism is "the holism of the right-making relation" (Dancy 2004, 79), while holism in the theory of reasons is "the holism of the favouring relation" (Dancy 2004, 79). In much of his work on moral particularism, Dancy seems to think of moral reasons as right-making factors (see, e.g., Dancy 1993, chaps. 4 & 6; Dancy 1999a, 59; Dancy 1999b, 25; Dancy 2000, 134-35), as does McNaughton (see McNaughton 1988, 190-94, 200-201; McNaughton 2000, 284-86;

distinguish these two views can cause confusion, including confusion about who is and who is not a moral holist, the question of whether we should include active background factors in moral and other reasons (as parts thereof) is simply beside the present point.[21]

Just how widespread moral atomism is may be an open question, but what is clear is that *prima facie* obligations are often understood as, or in terms of, invariant moral reasons or valence-invariant right-making factors.[22] What is also clear is that moral atomism is presupposed by many familiar, seemingly cogent moral arguments.[23] One such *atomistic argument* is James Rachels's argument against the view that killing someone is, in itself, morally worse than letting her die (1975, 228-29). Recall that Rachels asks us to consider "two cases . . . that are exactly alike except that one involves killing whereas the other involves letting someone die" (228), the case of Smith and the case of Jones (see page 5). And recall his claim that if these two cases are on a moral par, then "the bare difference

---

McNaughton and Rawling 2000, 257-58). That is, Dancy seems to think of the moral-favoring relation—the relation in which moral reasons stand to the actions for which they are reasons—and the right-making relation—the relation in which right-making factors stand to the rightness of the actions they make right—as the same relation, more or less. And in Dancy 2000, he argues from the holism of reasons generally to the holism of moral reasons, seemingly understood as right-making factors (131-37). But he now distinguishes the moral-favoring relation and the right-making relation, and argues from the holism of reasons generally to the holism of moral reasons, and then from the holism of moral reasons to the holism of right-making factors (Dancy 2004, 73-81; cf. Dancy 1993, 55).

[21] This certainly seems to be McNaughton and Rawling's view (2000, 271-72). For her part, Little does not even state her moral holism as a doctrine about moral reasons. Rather, what she says is that the factors in virtue of which acts have their moral properties do not "carry *their reason-giving force* atomistically" (emphasis altered) (2000, 280).

[22] See, e.g., Ewing 1959, 110-11; Frankena 1973, 26; Urmson 1975, 112-13; McNaughton 1996, 78-79, 81.

[23] This point has been made before, though in different terms—e.g., that they commit "the additive fallacy" (Kagan 1988). See also Dancy 1993, 60, 63-66, 88-90. Cf. Philips 1987.

between killing and letting die does not, in itself, make a moral difference" (229). Now

one might think that what Rachels means by this is simply that the bare difference in

question does not, in itself, make a moral difference *in the case of Smith*, or between it and

the case of Jones—although then his claim would be trivial, since if the two cases are on a

moral par there can't be any morally relevant difference between them. Or one might

think that he is merely arguing that this bear difference is *presumptively* irrelevant, morally

speaking, and challenging the advocate of its intrinsic moral relevance in other cases to

identify the relevant difference between his cases and others in which it is of intrinsic

moral relevance. But it is clear that what Rachels means is that the bare difference

between killing and letting die makes no moral difference, in itself, in *any* morally possible

case—i.e., that it is (in itself) invariantly morally irrelevant.

What make it clear that he intends this much stronger claim is his claim that the lack

of any morally relevant difference between the case of Smith and the case of Jones shows

that active euthanasia is not, in itself, worse than passive euthanasia *even though the two sorts*

of cases are nothing alike:

> Now, it may be pointed out, quite properly, that the cases of euthanasia
> with which doctors are concerned are not like this [the cases of Smith and
> Jones] *at all* [my emphasis]. They do not involve personal gain or the
> destruction of normal healthy children. Doctors are concerned only with
> cases in which the patient's life is of no further use to him, or in which the
> patient's life has become or will soon become a terrible burden.
> However, the point is the same in these cases: the bare difference
> between killing and letting die does not, in itself, make a moral difference.
> If a doctor lets a patient die, for humane reasons, he is in the same moral
> position as if he had given the patient a lethal injection for humane
> reasons. (229)

Moreover, what Rachels does here is *deduce* (*a*) 'That φ-ing would be a killing rather than a

letting-die is, in itself, morally indifferent *in any case*—i.e., that it is invariantly a morally

indifferent factor' from (*b*) 'That ɸ-ing would be a killing rather than a letting-die is morally indifferent *in the case of Smith*'. And from (*a*), he *deduces* (*c*) 'Active euthanasia is no worse in itself than passive euthanasia *in any case*—i.e., invariantly'. And this shows that he thinks there could be no relevant differences between his cases and others such that the difference between killing and letting die is of intrinsic moral relevance in the latter.

Rachels's argument presupposes moral atomism. That the cases of euthanasia with which doctors are concerned are nothing like his cases could be irrelevant, and his inferences could be valid, only if moral atomism were true. Moreover, his argument aims to establish a substantive moral claim—that there is no morally significant difference between killing and letting die—in order to establish another substantive moral claim—that active euthanasia is morally permissible. And it aims to do this without appealing to a background (normative) moral theory. That is, it aims to establish its conclusion as a stand-alone moral fact or principle with which any true or acceptable moral theory would have to be consistent. But for Rachels's argument to do that, he must be entitled *to assume* that the relevant factor's moral valence is invariant, and Rachels is not entitled to that assumption unless he is entitled to assume that right-making factors as such are valence-invariant. The only thing that might entitle Rachels to that assumption would be the truth of moral atomism: the conceptual impossibility of valence-variant right-making factors, or of variable intrinsic moral relevance. Otherwise, whether there are such factors and whether the fact that ɸ-ing would be killing rather than letting die is such a factor are clearly matters for substantive moral theory, matters on which different moral theories are free to disagree. Thus, Rachels's argument presupposes moral atomism.

Rachels's is far from the only atomistic argument. Peter Singer relies on a similar line of reasoning to argue that killing is (invariantly) no worse intrinsically than letting-die and, more generally, that there is no morally significant difference between acts and omissions (1979, 149-52). And this is a key premise in his defense of active euthanasia, and the sole premise from which he concludes both that pacifism's absolute condemnation of violence is indefensible (195-97) and that those of us in industrialized nations who have more than we need to maintain a basic level of welfare for ourselves not only have a moral duty to aid those in developing nations who do not have enough to maintain a basic level of welfare for themselves, but also that our duty to do so is on a moral par with our moral duty not to kill them (158-68).

Michael Tooley's (1972) much-anthologized defense of abortion and infanticide also relies heavily—and, if moral holism is true, fatally—on atomistic arguments. Because Tooley's positive argument rests on the controversial assumption that A (e.g., a fetus or an infant) has a right to x (e.g., life) only if A is capable of desiring x (Tooley 1972, 24), its cogency hangs on his criticisms of more intuitive alternative positions and principles. But his criticisms of the relevance of birth, viability, motility, and the potential for personhood to when a fetus or infant acquires a right to life all presuppose moral atomism (27-28, 31-32). For example, to support his contention that viability—the capacity for physiological independence—and its correlates are irrelevant to whether an organism has a right to life, he asks us to "consider a speculative case where a fetus is able to learn a language while in the womb" (Tooley 1972, 27), and he infers the irrelevance of viability and its correlates from the fact that such a fetus would have a right to life, even if it were not yet viable. Although it is not entirely clear that Tooley thinks this particular

inference is deductively valid, there is clearly the suggestion that neither viability, nor any of its correlates, nor any combination thereof, could make a morally relevant difference in actual cases unless it would make that same difference in speculative cases that are quite dissimilar, and hence the suggestion that any dissimilarities between such cases, whatever they may be, must be irrelevant to whether viability or its correlates (alone or in combination) make a morally relevant difference in actual cases. Otherwise, why think that such a dissimilar, speculative case is at all relevant to whether they make a morally relevant difference in actual cases? Moreover, it is not at all clear why a fetus's viability or its correlates (alone or in combination) could not make a difference (either as contributory factors or as active-background factors) to whether that fetus has a right to life just because other factors (e.g., the capacity for language acquisition and all that it requires) could also do so. In any case, Tooley's reliance on atomistic arguments renders his criticisms of more intuitive alternative positions and principles unsound and robs his positive argument of cogency.

In effect, each of these atomistic arguments follows Brandt in assuming that any evidence against the *universality* of a statement of the form 'there is a *prima facie* obligation to do A on account of F' is evidence against its *truth*. More generally, each takes evidence against the universality or invariance of a factor's moral relevance to be evidence against its intrinsic moral relevance. And if such "switching arguments" (Dancy 1981, 64-66) "are absolutely ubiquitous in ethics" (Hooker 2000, 6), then it is small wonder that moral atomism might be thought widespread. But whether or not their moral holism distinguishes moral particularists from most other moral philosophers, there are pro-

principle moral holists. Among these are Shelley Kagan and Thomas Hurka. Kagan is

clearly a moral holist, and Hurka is both a *value holist* and committed to moral holism.

3. The Additive Assumption and Moral Holism

Shelly Kagan rejects what he calls "the additive assumption" (1988, 19). This

assumption is the conjunction of three separate assumptions about how the various

factors that determine the overall moral status of particular acts (i.e., morally relevant

factors) combine and interact with one another:

1. The overall moral status of an act is the sum of the joint positive
   contribution made by the relevant positive factors and the joint
   negative contribution made by the relevant negative ones.

2. The joint positive contribution made by the relevant positive factors is
   (itself) the sum of the individual contributions made by those factors,
   and the joint negative contribution made by the relevant negative
   factors is (itself) the sum of the individual contributions made by
   those factors.

3. The nature of a given individual factor's contribution to the overall
   moral status of an act is determined independently of—i.e., is
   unaffected by—other factors.[24]

Together these three assumptions constitute a model of factor combination and

interaction for right-making and other morally relevant factors, namely, "the additive

model" (1988, 18). I will follow Kagan's usage and refer to the conjunction of these three

theses as "the additive assumption" (singular), but I need to refer separately to the

additive parts of the additive model—(1) and (2)—and the atomist part of it—(3). At the

---

[24] See Kagan 1988, 14-17.

risk of some confusion, I will refer to (1) and (2) collectively as *the additive assumptions* (plural). And I will refer to (3) as *the atomist assumption* (cf. Brink 1994b, 217 n. 5).

The latter, atomist assumption is even stronger than moral atomism.[25] It implies not only that if a factor f has a given moral weight and valence in one case, it must have that same moral *valence* in any case, but also that it must have that same moral *weight* in any case. In other words, it denies the weak dependent-contributions thesis, as well as the strong one. For this reason, we may distinguish a weaker atomist assumption. This *weak atomist assumption* denies only that a factor's moral valence may be affected by other factors, and so denies it may vary. It is thus equivalent to moral atomism. Kagan rejects not only the stronger atomist assumption, but also the weak atomist assumption (1988, 16 n. 7, 20, 24 & n. 11). This makes him a moral holist.

Kagan argues that the additive model is too simple to capture numerous plausible views, including the view that "the suffering of the guilty simply does not count as much as the suffering of the innocent" (1988, 20). He asks us to imagine that we can help either Trixie or Fritz, both of whom are suffering because Trixie's scheme to harm Fritz

---

[25] The additive assumption entails what Kagan calls "the ubiquity thesis," the thesis that if a factor makes a moral difference in any case then it must make a moral difference in every case (1988, 12-13). At first glance, the ubiquity thesis looks like moral atomism, but it is weaker. It implies only that if a factor has a moral valence in one case, it must have a moral valence in any case. It does not imply that if x has a moral valence in one case, it must have *the same valence* in any case. That is, it allows that a moral factor may be a positive—or right-making—factor in one case and a negative—or wrong-making—factor in another, which is something moral atomism denies is possible. As Kagan puts it: "The mere satisfaction of the ubiquity thesis does not itself guarantee that a given value of a factor will always make a contribution of the same size (nor even that the direction of the contribution will be constant—i.e., always positive or negative)" (Kagan 1988, 16 n. 7). The ubiquity thesis also allows "Janus-like" (Dancy 1993, 62) moral reasons. A Janus-like moral reason is a factor that contributes both positively and negatively to the overall moral status of the same (token) action. Kagan (1988, 29) claims we cannot assume that there are no such factors.

went awry. On the additive model, Trixie's suffering, Fritz's suffering, Trixie's guilt, and

Fritz's innocence are four mutually independent contributing factors whose sum

determines which of them we should help.[26] It can "point to [Trixie's] guilt as a reason

for not helping Trixie, but it is forced to maintain that the magnitude of her suffering

provides just as strong a reason to aid her as if she were innocent" (1988, 20). To capture

the view that Trixie's suffering counts less because of her guilt, we need a different model

of moral factor combination and interaction.[27]

Now, Kagan does not claim that this view is true, or that any other view the additive

model cannot capture is true. But he does claim that because the additive model cannot

capture this and other plausible views, those who hold such views "will find it plausible to

reject the additive assumption" (1988, 20). Kagan suggests that one might do this by

adopting either of two seemingly alternative models of factor combination and

interaction. The first of these models gives up the assumption that the contributions

made by individual factors must be summed—i.e., the additive assumptions. Applying

this model to the case of Trixie and Fritz, one might treat Fritz's innocence as a multiplier

of one ($x = 1$), so that his suffering (y) has its full potential weight ($S = xy$), but treat

Trixie's guilt as a multiplier of less than one ($x < 1$), so that her suffering has less than its

---

[26] Kagan treats innocence/guilt as two possible values of a single factor, culpability. He suggests, however, that at least in many cases, the choice of whether to speak of an always-present factor, like culpability, or a one-valued factor that is not always present, like guilt, will be merely "a matter of terminological convenience" (1988, 6 n. 3).

[27] Might one take there to be two, rather than four, mutually independent factors here (*viz.*, Trixie's guilty suffering and Fritz's innocent suffering) whose sum determines whom we should help? Perhaps, but that wouldn't change the fact that the additive model cannot capture the fact that Trixie's suffering counts less because of her guilt. To do that, we need to treat her guilt and her suffering as distinct factors.

full potential weight ($S < y$). And one could, Kagan says, allow for the deserved suffering of the wicked by allowing the value of the guilt multiplier to range below zero ($S < 0$), so that such suffering could have negative weight.

The second model gives up the assumption that the nature of a given individual factor's contribution (i.e., its weight and valence) is determined independently of other factors. That is, it gives up the atomist assumption. Applying this model to the case, one might treat Trixie's guilt as a factor that affects the nature of the contribution made by Trixie's suffering, without itself contributing one way or the other to the rightness or wrongness of helping her (or of helping Fritz). On both models, there are four morally relevant factors: Fritz's innocence, Fritz's suffering, Trixie's guilt, and Trixie's suffering. On the first model, each of these factors makes a separate and independent contribution to the moral status of helping Fritz, but the moral status of doing so is not the sum of these contributions. On the latter model, only two of them—Fritz's suffering and Trixie's suffering—contribute to the moral status of helping Fritz, and their contributions are dependent on the remaining two factors—Fritz's innocence and Trixie's guilt.

These two models appear to be quite different. In particular, on the first model the nature of a given individual factor's contribution to the overall moral status of an act is determined independently of other factors, while on the second model this is not the case. Thus, the first looks like a model of an atomistic view, albeit a non-additive one; and the second looks like a model of a holistic view, albeit one that (unlike moral holism) is committed to the additive assumptions. Yet, despite how different these models appear, it is not at all clear that there is any substantive difference between them. This is because if one describes the latter, seemingly holistic model mathematically (i.e., in

mathematical terms), the description is identical to that of the former, seemingly atomistic model.

Indeed, the former view looks to be nothing other than a mathematical model of moral holism. If one were to model a right-making factor whose weight or valence could vary—i.e., a variant moral reason—in mathematical terms, one would have to model it as a value fixed independently of other values. One would have to model it as the value of the contribution the factor would make in the absence of any active background factors that might alter its normal weight or valence—i.e., its normal value. If one were to model an active background factor mathematically, one would have to model it as a multiplier (of the normal value of the contributing factor it affects). And, if one were to model the actual contribution a variant moral reason made in a particular case in the language of mathematics, one would have to model it as the product of a value fixed independently of other values and one or more multipliers, for one would have to model it as the product of the factor's normal value and the values of any active background factors. In other words, one would have to model it as the product (rather than the sum) of mutually independent values. Thus, a mathematical model of moral holism would be identical to the seemingly atomistic model of factor combination and interaction that Kagan describes. For this reason, I will treat these two models as equivalent.

For his part, Kagan is more or less indifferent between the two models. He stresses that he "would not want to put much weight on th[e] mathematical analogy" (1988, 19) and that it too may well be inadequate (1988, 20). He also acknowledges one might find

the other model more natural (1988, 20-21 n. 8),[28] makes use of both models (see, e.g.,

1988, 24), and even has a tendency to run them together; as when he says that to those

who believe the suffering of the guilty does not count as much as that of the innocent "it

will seem more appropriate to view the factor of innocence/guilt not as making *an*

independent contribution, but rather as *a multiplier*" (my italics) (1988, 20). His aim is simply

to unmask the additive assumption for what (he thinks) it is: a controversial, substantive

claim that is inconsistent with most moral views and does a poor job of capturing our

intuitions about how morally relevant factors combine and interact with one another in

particular cases. Moreover, in more recent work, he seems to have given up on the

mathematical analogy and model in favor of the terminology of moral holism (Kagan

1998, 307).

Now I am not claiming that Kagan believes there *are* variant moral reasons—i.e.,

that he is an ontological holist. He argues against the additive *assumption*, not the additive

model. That is, he argues only that we cannot *assume* the additive model is an accurate

---

[28] As between the two models, the latter does an intuitively better job of accommodating many, if not all, of the intuitions Kagan wants to capture. Consider, for example, the case of Trixie and Fritz. Both Trixie and Fritz are suffering because Trixie's scheme to harm Fritz went awry. The non-additive, seemingly atomist model might treat Fritz's innocence as a multiplier of one, so that his suffering has its full potential weight, and treat Trixie's guilt as a multiplier of less than one (and perhaps less than zero), so that her suffering has less than its full potential weight. On this model, Trixie's suffering still has the moral significance it would normally have; but, because of her guilt, at most only a fraction of that significance contributes to the rightness of helping her. This is counterintuitive. And it seems to suggest something else counterintuitive, that her suffering should move us to help her just as much as it normally would but that we should nonetheless discount it in our deliberations. Intuitively, her suffering is not as significant as it would normally be: it neither counts as much as it normally would nor should it move us as much to help her. The holistic model can capture these intuitions, for it can say that Trixie's guilt affects the moral significance of Trixie's suffering itself, so that her suffering lacks the moral significance it would normally have. Thus, of the two models, the holistic one is preferable, at least in this respect.

model of how morally relevant factors combine and interact (and that arguments that do so are fallacious—that they commit "the additive fallacy"). And one can accept *that* and yet believe the additive model is, in fact, an accurate model of how such factors combine and interact. But I am claiming he is a moral holist. For he allows that right-making factors as such are not valence-invariant, and he regards the question of whether there are variant moral reasons as a substantive question to which different moral theories may give different answers:

> Once one entertains a picture of normative ethics with several morally relevant factors, the structural possibilities get quite complex (for example, factors can override, outweigh, or undermine one another, or interact in a variety of other ways). (Kagan 1998, 307)

Moreover, inasmuch as the two alternative models of factor combination and interaction Kagan considers are equivalent, his argument against the additive assumption is an argument for moral holism.

4. Value Holism

Kagan remarks (in a footnote) that "Moore's defense, in *Principia Ethica*, of an organic theory of the good can be seen as a rejection of the additive assumption in a particular area of value theory" (1988, 23 n. 10). Moore's principle of organic unities states:

> The value of a whole must not be assumed to be the same as the sum of the value of its parts (Moore 1903, 28).

The kind of whole Moore has in mind here is one composed of goods, evils, and indifferents—that is, things good, bad, or indifferent *in themselves*, or *intrinsically*. Hurka glosses Moore's principle as:

> If two states are brought together so a certain relation holds between them, the resulting whole may have either more or less intrinsic value than the states would have if they existed alone (Hurka 1998, 299).

And he claims that it is open to two interpretations, Moore's own interpretation and a different "conditionality" (1998, 299, 303) or "variability interpretation" (2003, 604-7). On the former, goods, evils, and indifferents have the same intrinsic value within wholes as they do outside of them, but the wholes composed of them may have intrinsic values of their own *as wholes*. Thus, the intrinsic value of a whole is not equal to the sum of the values of its parts, but rather to the sum of the values of its parts *and* the whole's own (intrinsic) value as a whole. On the latter interpretation, goods, evils, and indifferents may have values within wholes that differ from the values they have outside of them. Goods may be less good, indifferent, or even evil within wholes. Evils may be less evil, indifferent, or even good within wholes. And indifferents may be good or evil within wholes. Thus, the intrinsic value of a whole may be either more or less than the sum of the values its parts would have on their own.

The view of intrinsic value that results from the former interpretation (Moore's own) is an atomistic one. Although Hurka refers to Moore's own interpretation of the principle as "the holistic interpretation," from the perspective of moral holism Moore's view of intrinsic value is *atomistic*, because it holds that if a thing has a given intrinsic value in one case, it has that same intrinsic value in any case. But the view of intrinsic value that results from the latter, conditionality or variability interpretation is an analogue in the

theory of intrinsic value of moral holism (and of Kagan's seemingly atomistic, non-additive model of factor combination and interaction.)[29]

On this conditionality, or variability, view of intrinsic value, the intrinsic value of a thing—its value in itself, apart from its consequence—may be affected by things other than itself or by the relations in which it stands to other things.[30] Things good and evil in themselves when alone or under normal conditions may be neutral in themselves under certain conditions. For example, pleasure may be good in itself both when it is experienced alone and under normal conditions, yet lack value when contrary to desert (cf. Feldman 1995, 577) or when taken in another's pain. Things indifferent in themselves when alone or under normal conditions may be good or evil in themselves under certain conditions. Thus, knowledge and other perfectionist goods may be good in themselves but only when accompanied by pleasure in them (cf. Frankena 1973, 90-91).[31] Or things good or evil in themselves when alone or under normal conditions may be "transvalued"—that is, something that is good in itself when alone or under normal

---

[29] Although it may appear as if Moore's own interpretation of his principle is an analogue of Kagan's model, it is not. Kagan's model represents the moral status of an act as a non-additive function of values, none of which represents the act itself. But Moore's view maintains that the value of a whole is equal to the sum of the values of its parts and of itself as a whole. Thus, if one were to model his view mathematically, one would have to include not just values describing the parts of wholes, but additional values describing the values of the wholes as wholes; and the function governing the overall intrinsic value of wholes would be an additive one.

[30] Hurka discusses this view in a number of places. See esp. Hurka 1998; Hurka 2001, 181-90; Hurka 2003, 604-7.

[31] Depending on how "normal conditions" is understood, a thing could be indifferent in itself when alone but good or evil under normal conditions. Suppose knowledge is usually accompanied by pleasure in it, but is indifferent when alone. It would then be indifferent when alone but good under *statistically* normal conditions. If we do understand normal conditions statistically, then active background factors may be present in normal conditions. In particular, factors whose absence may affect a thing's value may be present.

conditions may be evil in itself under certain conditions, and vice versa. Thus, pain may

be good, rather than bad, when it is deserved (cf. Feldman 1995, 579) or when felt as a

response to some evil. Indeed, if we think of individual (intrinsic) goods and evils as

factors that contribute to the (overall) intrinsic value of states of affairs, we can state this

view in terms that parallel those of moral holism. In particular, we can state a strong

dependent-contributions thesis for *good-making factors*:

> The magnitude and valence of a factor f's contribution to the (overall)
> intrinsic value of state of affairs s may be affected by factors other than f
> itself.

For convenience, I will refer to this view as *value holism*.[32]

Hurka is a value holist. Again, this is not to say he thinks that there are variant good-

making factors, or values. But he clearly regards claims that there are such things as both

coherent and plausible, and he treats them as substantive claims that must be met on

substantive grounds, not as conceptual impossibilities, or as claims to be settled prior to

substantive inquiry, reflection, and theorization.[33] Thus, he allows that the magnitude and

valence of a good-making factor f in circumstance c may vary as the effect of factors other

than f itself in just the sense that moral holists believe that the weight and valence of a

right-making factor f in circumstance c may vary as the effect of factors other than f itself.

Given this, one might simply infer that he is also a moral holist. After all, there is no

obvious reason why someone who accepted value holism would deny moral holism, and

---

[32] Not surprisingly, Dancy both endorses this view and calls it both "holism in the theory of value" and "particularism in the theory of value" (Dancy 2000, 137-41; Dancy 2003). Feldman and Ross both recognize variant (intrinsic) goods (Feldman 1995, 576-80; Ross 1930, 136-38).

[33] See, e.g., Hurka 1998; Hurka 2001, 144-52.

one might well expect that if one were true, the other would be as well. But it is worth

showing that Hurka is, in fact, committed to moral holism, and why.

Hurka allows that the fact that ɸ-ing would promote x is a right-making factor (it

contributes to the rightness of ɸ-ing) if x is good but is a wrong-making factor if x is

evil.[34] And it follows from this that if some x (e.g., pleasure) may be good in some

circumstances and evil in others, the fact that ɸ-ing will promote that x is a variant moral

reason. It follows that Hurka is a moral holist, or at the very least committed to moral

holism. Indeed, any moral theorist who both accepts this principle—including any

consequentialist—and is a value holist is *ipso facto* committed to moral holism.[35]

---

[34] See Hurka 2001, 4-5, 7, 12; see also Hurka 1996, 557 n. 1.

[35] One might object that since consequentialist moral theories typlically identify right
actions by reference to the intrinsic value of states of affairs, they can only recognize a single,
invariant right-making factor (e.g., that ɸ-ing would maximize aggregate intrinsic value). But this
rests on a confusion between two different understandings of what a right-making factor is. It
supposes a right-making factor is that property shared by all and only right acts in virtue of which
they are right. But the relevant understanding of a right-making factor is that of a factor that
makes a positive contribution to the rightness of particular acts, or counts morally in favor of
performing them. A consequentialist moral theory can recognize a plurality of such factors, and it
can recognize variant ones.

Section I of this chapter, in part, has been submitted for publication. I am the sole

author.

CHAPTER 4

A DEFENSE OF MORAL HOLISM

One could take any number of views with respect to cases like Dancy's worm case, and with respect to the fact that φ-ing would give someone pleasure. Some will say just what Dancy and many other moral holists say about the latter, that it is a variant moral reason. Others will say that it is an invariant one, that it makes φ-ing *pro tanto* right in any morally possible case, even one in which φ-ing consists in harming or even killing another sentient being. Still others will say that the mere fact that φ-ing would give someone pleasure is morally indifferent. And some (but not all) of those who say the latter will also say that the fact that φ-ing would give someone *innocent* or *non-sadistic* pleasure is an invariant moral reason to φ. Hooker, for example, wants to say that the fact that φ-ing would produce *non-sadistic* pleasure is an invariant moral reason to φ—though he does not "pretend to confidence" that the fact that φ-ing would produce *sadistic* pleasure is an invariant moral reason *not* to φ (Hooker 2000, 8). (It seems not to occur to Hooker that there is a difference between his own view and the view that the moral valence of the fact that φ-ing would produce *pleasure* depends on what the object of the pleasure is, nor does he offer any reason to think that his view is either correct or preferable to the latter one.) But others may want to argue that other factors—considerations of desert, for example—can alter the moral valence of the fact that φ-ing would produce non-sadistic pleasure.

Much the same can be said about the case of Smith, and about the fact that ϕ-ing would be killing rather than letting-die. Some will say that the fact that ϕ-ing would be killing rather than letting die is a variant moral reason (cf. Kagan 1988, 18-19). But others will say, with Rachels, that it invariantly makes ϕ-ing *pro tanto* wrong, even though its moral relevance in some cases, including the case of Smith, may fall below "the threshold of ordinary discernability in the context of decision-making" (Audi 1998, 38). Similarly, some will say that someone's guilt (e.g., Trixie's) or considerations of desert can affect the moral valence of the fact that ϕ-ing would relieve someone's suffering.

Indeed, intuitions about these and other cases and factors will likely differ considerably, as will the accounts of them that different moral theories offer. And in each case, it is a substantive question which if any of the competing intuitions or accounts is the correct one. Or, rather, so says the moral holist. For the moral atomist says something quite different. The moral atomist claims that we need not consider holistic accounts of such cases and factors. Indeed, the atomist claims that we *cannot* consider them. For moral atomism rules out variant moral reasons, and hence holistic accounts of such cases, as conceptual impossibilities.[1] And this is the basis for what I call *the argument from examples*.

---

[1] Or rather, it holds that if a factor is a variant moral reason, then it must be a *derivative* right-making factor (i.e., a factor whose moral relevance derives solely from the relations in which it stands to one or more other right-making factors). For it holds it to be a conceptual truth that non-derivative right-making factors are valence-invariant (see Chapter 3).

1. The Argument from Examples

Consider the following two cases offered by Dancy:

D1: "I borrow a book from you, and then discover that you have stolen it from the library. Normally the fact that I have borrowed a book from you would be a reason to return it to you, but in this situation it is not. It isn't that I have *some* reason to return it to you and more reason to put it back in the library. I have no reason at all to return it to you" (Dancy 1993, 60).

D2: "[Imagine] a family game called 'Contraband', in which the players are smugglers trying to get contraband material past a Customs Officer. The game requires them to lie; if one doesn't do plenty of lying, it spoils the game. That an action is a lie is commonly a reason not to do it; here it is reason in favour" (Dancy 1993, 60-61).

And consider this pair of contrasting cases offered by David McNaughton:

M1: "I take my nephews and nieces to the circus for a treat. They enjoy it. I have done the right thing" (McNaughton 1988, 193).

M2: "A government is considering reintroducing hanging, drawing and quartering in public for terrorist murders. If reactions to public hangings in the past are anything to go by a lot of people may enjoy the spectacle. . . . Is the fact that people would enjoy it here a reason for its being right? It would be perfectly possible to take just the opposite view. The fact that spectators might get a sadistic thrill from the brutal spectacle could be thought to constitute an objection to reintroduction" (McNaughton 1988, 193; cf. Dancy 1993, 61).

One might take either of these or any of the other cases offered by holists to show that a particular factor's moral valence is, in fact, variable; and if one factor's moral valence is variable, then moral holism is true. But one might not share the particular intuitions appealed to. Moreover, even if one does, one might—quite rightly—take them to be revisable. So just what do such putative counterexamples to moral atomism show?

Perhaps they do not show that moral atomism is false, and that moral holism is true. And it might yet be that, on the best conception of what such factors are, right-making factors must be valence-invariant. But such examples do show at least this much—that moral atomism is a substantive view, and thus that it may not be assumed without argument (cf. Kagan 1988, 13, 18).[2] Moreover, they are sufficient to raise a strong presumption in favor of moral holism.[3]

Consider Dancy's examples. Dancy plausibly claims that each of his examples is one "in which the most *natural* thing to say is that a consideration is functioning in one case as a reason in favour of doing the action, but in another is either no reason at all, or even a reason against" (1993, 60).[4] In D1, the fact that Dancy borrowed a book from you appears not to be a reason (moral or otherwise) for him to return that book to you, even though in many other cases it would appear to be a moral reason for him to return it to you (a right-making factor). In D2, the fact that if Dancy were to φ he would be lying appears to be a reason for him to φ, even though in many other cases it would be a reason for him not to φ (a wrong-making factor). If moral holism is false, all this is conceptually impossible.

---

[2] For other examples, see Dancy 1993, 60-62; McNaughton 1988, 192-94. The examples in Kagan 1988 (18-21) may also be taken to support moral holism.

[3] They are also, I think, sufficient to raise a (weaker) presumption that weaker forms of atomism are false, but my argument does not require this.

[4] Dancy suggests he originally intended his examples to support holism in the theory of reasons (or the holism of the favoring relation), rather than moral holism (or the holism of the right-making relation). But, as he notes, they are easily adapted to the cause of the latter (Dancy 2004, 79-80), and that is how I use them.

Consider M1 and M2. McNaughton's intuition about M1 is that it was because his nieces and nephews enjoyed it—because he succeeded in giving them pleasure—that it was right for him to take them to the circus. A moral atomist who accepts this must also accept that the enjoyment or pleasure that many people would get from the spectacle of public hangings is a moral reason for reintroducing them, a factor that makes doing so *pro tanto* right. But to McNaughton this seems wrong. His intuition is that any pleasure that would be derived from the spectacle of public hangings is not merely irrelevant to whether they ought to be reintroduced, but a moral reason *not* to reintroduce them—it would make doing so *pro tanto* wrong. And, as he implies, this intuition about M2 is both plausible and consistent with his equally plausible intuition about M1. Why think, as the moral atomist does, that McNaughton's intuitions are inconsistent?

Kagan, who is more explicit about the nature of this argument, rightly claims that such examples show the additive assumption and its various parts—including the weak version of the atomist assumption—to be substantive theses that, as such, may not be assumed without argument. (He also claims they show them to be controversial theses.) Here is one pair of (contrasting) cases he offers:

> **K1:** In a legitimate (i.e., morally permissible) case of self-defensive, I push an aggressor into a pit, expecting the fall to kill him.

> **K2:** In a legitimate case of self-defense, I refrain from warning an aggressor about the pit into which he is about to fall, and I expect the fall to kill him.

Kagan claims not only that "most of us will certainly want to claim that the moral status of these cases is the same," but also that "most of us will want to maintain that the distinction between doing harm and allowing harm makes *no* difference in such legitimate

cases of self-defense" (Kagan 1988, 18). That is, Kagan claims most of us will share not just his intuition that it is morally permissible to do harm in the first case and allow it in the second but also his intuition that there is no morally significant difference whatsoever between these two cases: that they are morally equivalent. Many, of course, have the opposite intuition in many cases; they intuitively find doing harm worse than allowing it, even when it is nevertheless permissible to do harm. And it might be that some, or even most, won't share Kagan's intuitions about these cases. But what Kagan rightly points out is that one might well hold both sets of intuitions, and that it is perfectly coherent to think that the fact that one would be doing harm rather than allowing it is typically one of great moral significance but is of no moral significance whatsoever in cases of legitimate self-defense. As he says, "the crucial point to see is that even the *advocate* of the do/allow distinction may well want to make this claim, and surely it is coherent for her to do so" (Kagan 1988, 18).[5]

The crucial claim here is that the intuitions to which Dancy, McNaughton, and Kagan appeal are plausible, consistent, and reflect (or at least could reflect) coherent views about (*a*) what factors are moral reasons—right- and wrong-making factors—and (*b*) the conditions under which they are such. This is the crucial claim, because to accommodate *the mere conceptual possibility* that these intuitions are correct (their mere conceptual coherence), we must allow that moral holism is true (cf. Kagan 1988, 19). If moral holism is false, then not only must the fact that A borrowed a book from B make

---

[5] Indeed, he claims "most of us want to claim" not just that this is a coherent thought, but "that it is true" (Kagan 1988, 19).

A's returning it to B *pro tanto* right in either *every* possible case or *no* possible case, but it is also *inconceivable* that things are otherwise. And if moral holism is false, then not only must the do/allow distinction (and the killing/letting-die) distinction be morally significant in either every possible case or no possible case, but it is also *inconceivable* that things are otherwise. Likewise, if moral holism is false, then not only must the fact that φ-ing would promote pleasure make φ-ing *pro tanto* right in every possible case if it does so in any possible case, but it is also *inconceivable* that things are otherwise. (So too for the fact that φ-ing would relieve someone's suffering.) For moral atomism says that 'what is a right-making factor in one case may *not* have a different moral valence in any case' is a conceptual truth.

Granted, moral atomism derives some plausibility from the fact that it *seems* inconsistent to say something like, "I grant you that the fact that your nieces and nephews would enjoy the circus is a reason to take them to it. But I deny that the fact that many would enjoy the spectacle of public executions is a reason to reintroduce them. Indeed, I think it is a reason not to do so." Someone who makes such a claim might be asked to explain how it could be true. But a plausible explanation of how it could be true would rebut any charge that it is inconsistent (cf. Dancy 1993, 63-64). And what the moral holist claims is that there are plausible explanations available for why the moral valences of the factors in their examples might vary as they seem to do—explanations that appeal to other factors and their effects on the moral valences of the right-making factors in question. Dancy claims that the reason he would have no reason to return the book to you in D1 is that that you stole the book from the library. And he claims the fact that an action is a lie is a reason to do it when playing the game Contraband (as in D2) because

Contraband requires plenty of lying and will be spoiled if one does not lie. Similarly, McNaughton claims it is because the source of the spectator's enjoyment in M2 would be the brutality of the spectacle of public hangings that their enjoyment is a reason not to reintroduce such hangings. And Kagan implies that one may appeal to the fact that a case is one of morally permissible self-defense to explain why the fact that one would be doing rather than allowing harm makes no moral difference to φ-ing therein, or to the fact that the suffering in question is deserved to explain why the fact that φ-ing would relieve suffering makes φ-ing *pro tanto* wrong in that case. Each of these is a plausible explanation of why the moral valences of these factors vary (or might vary) in these cases, and so rebut any charge of inconsistency.

Russ Shafer-Landau questions the probative value of Dancy's examples. He claims that examples, like D1, in which a factor appears to lack its normal valence cannot show that it does lack that valence because "we lack an account of how to decide in a given case whether a property's omnipresent relevance is being overridden, or whether the property isn't always relevant" (Shafer-Landau 1997, 590). But insofar as this remark concedes that a factor may have a positive or negative moral valence in one case but no such valence in another, it concedes the truth of moral holism. Moreover, which way does Shafer-Landau's claim cut? Granted, one might insist that the moral relevance of a factor can fall below "the threshold of ordinary discernability in the context of decision-making" (Audi 1998, 38) without falling to zero. But what reason is there not to take an example like D1 at face value, especially when there is a plausible explanation available of why the relevant factor lacks its normal valence?

Shafer-Landau also presses the following objection against examples, like D2, in which the moral valence of a factor appears to reverse:

> Dancy's examples may tempt one to the view that sees a property's value [its moral valence] as entirely reversible depending on context only because he has focused on candidate properties [factors] that are very generally described. But once we narrow the scope of the grounding properties [the right-making factors], it is not very difficult to identify some that are universally and uniformly relevant [valence-invariant]. Taking pleasure at another's pain, or another's undeserved pain, or from the performance of a wrong action is always presumptively wrong. (Shafer-Landau 1997, 590)

This is a non sequitur. Moral holism claims that distinct factors can interact such that one reverses or otherwise alters the moral valence of the other: that, for example, the fact that the pleasure would be taken in the pain of another reverses the moral valence of the fact that φ-ing would give someone pleasure, just as the batting of a pitched baseball reverses its direction of travel. It would be absurd to say that the latter example may tempt one to the view that sees a baseball's direction of travel as being entirely reversible depending on context only because it is very generally described, and that once we narrow the scope of the relevant description, we see that batted baseballs always travel in one direction. A moral holist could perfectly well say that it is always presumptively wrong to give someone sadistic pleasure *because* the fact that the pleasure would be taken in the pain of another is capable of reversing the moral valence of the fact that φ-ing would give someone pleasure, thus making it *pro tanto* wrong to φ. The moral holist's response to Shafer-Landau is that he is narrowing the scope of the relevant description to the point that it excludes explanatory interactions, and that in doing so, he leaves it obscure why taking pleasure in another's pain is presumptively wrong but taking pleasure in, say, eating ice cream is not. His preferred descriptions are too narrow; they describe

the *explananda*, the effects of interaction, but exclude the *explanantia*, the interactions themselves.

We should conclude, then, that the argument from examples not only shows moral atomism to be a substantive view, but that it also establishes a strong presumption in favor of moral holism and shifts the burden of argument to the moral atomist.[6] Granted, our holistic intuitions would have to yield to a compelling atomistic account of what such factors are and of how they must behave, just as they would have to yield to a compelling account of how they *do* behave (cf. Dancy 1993, 62, 105). But it would count against either sort of theory that it discounted all of our holistic intuitions, or explained them all away as mistaken. Moreover, the moral atomist is committed to saying that holistic intuitions about such cases *must* be mistaken, that they are *inconsistent*. Indeed, the moral atomist claims that variant moral reasons are *inconceivable*. And these are *exceptionally* strong claims, on a par with: 'It is inconceivable that whether striking a match will cause it to light could depend on whether the match is wet or dry, on the presence or absence of oxygen, or on any other factor; for it is a conceptual truth that either striking a match must always cause that match to light or it must never do so'. Indeed, it is hard to imagine just what an argument for such a view would look like.

---

[6] If we distinguish between moral holism and particularism in theory of moral reasons, it is also, *mutatis mutandis*, sufficient to establish a strong presumption in favor of the latter, and for holism in the theory of reasons, of which it is a special case (cf. Dancy 2004, 79-80). Intuitively, moral and other reasons behave holistically, and so examples like those offered by Dancy, McNaughton, and Kagan not only show that atomism about moral and other reasons is a substantive view, they also raise a strong presumption against it.

2. Holism and Universalizability

Crisp objects that holism in the theory of reasons implies that moral judgments are not universalizable. Specifically, he claims moral holism is inconsistent with the thesis that "if some action is judged to be right, then any relevantly similar action must likewise be judged to be right" (Crisp 2000, 40), which he claims "seems plausible" (42). Although this is an objection to holism in the theory of reasons, rather than moral holism, it warrants answering here. For if those who conceive of right-making factors as (deontic) reasons are right to do so, then moral holism entails holism in the theory of reasons.

Crisp's thesis is vague; and on some interpretations it is clearly false, while on others it is trivial. I need not judge that some action is right even though I judge that it is relevantly similar to another action I judged to be right in the past. I am not constrained by my past judgments in this way, because I may come to see my past judgments as mistaken. Moreover, I may mistakenly judge that two actions are not relevantly similar; in which case I may consistently judge that one is right and the other not. Of course, if I simultaneously judge that two actions are relevantly similar and that one is right, then I must—on pain of contradiction—judge that the other is right. But moral holism is perfectly consistent with this trivial thesis. Thus, Crisp's objection is not well taken, and it has no force against moral holism, even if the latter does entail holism in the theory of reasons.

Could some version of the thesis that moral judgments are universalizable ground an objection to moral holism? I do not see how. The moral holist will allow, of course, that it would be inconsistent to deny that two actions one judges to be relevantly similar

are different. That much is trivial. And if a moral holist accepts that moral properties supervene on non-moral ones, then she will hold that one is *mistaken* if one judges that two actions that are identical in their non-moral properties differ in their moral ones. But the moral holist will say, quite reasonably, that any claim that two actions are relevantly similar or dissimilar is a substantive moral claim. And thus, she will reject any attempt to establish a formal constraint on relevant similarity as question begging.[7]

3. The Aim of Ethical Theory

Crisp also insists that the aim of ethical theory "is to locate reason-giving properties of which we can safely assume the 'additive assumption' to hold true" (Crisp 2000, 35 n. 38; cf. Kagan 1988). Since the 'additive assumption' includes, as parts, both the denial of the strong dependent-contributions thesis and the denial of the weak dependent-contributions thesis, such properties would be ones with invariant moral *weight* as well as invariant moral *valence*. Yet Crisp does not deny the weak dependent-contributions thesis; rather he claims that we should speak of complexes of contributory factors and other factors relevant to their weights as having weight "*as wholes*, and of their constituents as individually having certain weights *in the absence of* the other constituents" (Crisp 2000, 47 n. 73). This suggests that we might think of such complexes as weight- and valence-invariant *second-order* right-making factors. But why should we take the aim of ethical theory to be an atomistic moral theory, let alone a *fully* atomistic one, in the first place?

---

[7] For more on holism and universalizability, see Dancy 1993, 55-58, 79-81, 88-90, 258; Dancy 1999b; Sinnott-Armstrong 1999.

Crisp takes Newtonian mechanics as his model (Crisp 2000, 35), and Newtonian

mechanics is atomistic. It says that the individual forces acting on an object make

valence-invariant contributions to the total force acting on that object, which in turn

determines the direction and rate of its acceleration. But the suggestion that Newtonian

mechanics can serve as a model for moral theory is dubious. Unlike Newtonian

mechanics, chemistry is *holistic*:

> When two forces in mechanics are present together, each retains its
> original capacity. They operate side by side, independently of one
> another. The resulting effect is a pure combination of the effect that each
> is trying to produce by itself. The law of vector addition gives precise
> content to this idea of pure combination. In chemistry, things are
> different. The acid and the base neutralize one another. Each destroys
> the chemical powers of the other, and the peculiar chemical effects of
> both are eliminated. This is not like the stationary particle, held in place
> by the tug of forces in opposite directions. When an acid and a base mix,
> their effects do not combine: neither can operate to produce any effects
> at all. (Cartwright 1989, 163)

It might just as well be that morality is more like chemistry in this respect than it is like

Newtonian mechanics (cf. Kagan 1988, 13).[8]

Indeed, the more plentiful and plausible the examples offered by moral holists are

(and the more they feature factors whose moral valences others have taken to be

invariant), the more they suggest that morality *is* more like chemistry in this respect—that

it *is* holistic. Moreover, there is independent reason to reject Newtonian mechanics as a

---

[8] Why doesn't the vector addition model not also give precise content to the way in which
acids and bases combine? First, the relevant claim is that it doesn't give precise content to the
way in which *their effects* combine, not the way in which they combine. Second, when acids and
bases combine, they cease to exist, and hence to have any effects. Third, when acids and bases
combine, their chemical *powers* don't. Instead, they *interact*. And in that interaction, they too are
destroyed. The vector addition model cannot represent the underlying metaphysics here, and
hence it cannot give precise content to either the way in which acids and bases combine or the
way in which their powers interact.

model for moral theory: the result of interactions between right-making factors is not a pure combination of the effect that each is trying to produce by itself. Even if something like a vector addition model might be useful *for determining* (i.e., for figuring out) what one's obligation overall is when one's (*prima facie*) obligations conflict, when one's obligation overall is to act just as one of one's conflicting obligations would require in the absence of the others, this result will be neither *determined by* nor *equal to* a vector sum of their individual weights.[9] If morality doesn't fit the model of Newtonian mechanics in this respect, what reason is there to think that Newtonian mechanics can serve as a model for moral theory?[10]

The suggestion that an atomistic moral theory ought to be the aim of moral theory is just as dubious, if not more so. We know that mechanical forces are atomistic (if indeed we know this) by doing physics, and we know that chemical powers are not (if indeed we know this) by doing chemistry. Presumably, then, we learn whether right-making factors, forces, or powers are atomistic or holistic by doing normative moral theory, not through conceptual analysis or metaphysical speculation that is, relative to such theorizing, *a priori*. To know what morality is like, we must do *normative* moral theory. And therein lies the rub. How could we learn whether morality is holistic rather than atomistic by doing normative moral theory if we take the aim of moral theory to be

---

[9] Cf. Kagan 1988, esp. sec. III & sec. IV; Dancy 1993, 101-4; Brink 1994b, sec. 1 & nn. 5 & 13.

[10] Cartwright claims "classical mechanics may well be the only discipline where a general law of action [*viz.*, the law of vector addition] is always available" [1983, 65]. This suggests that classical mechanics may be *unique* in being atomistic, or at least in having forces that combine additively—if those are even two different things, rather than two different ways of saying the same thing.

an atomistic moral theory? It would be tantamount to doing chemistry all the while assuming that chemical powers must combine in just the way that Newtonian mechanics says that forces do, regardless of any evidence to the contrary.

The suggestion that we can secure invariance by appealing to second-order right-making factors of the sort that Crisp envisions is also dubious. To see why, consider the attempts of Hooker and others to offer counterexamples to the claim that there are no invariant moral reasons.[11] Obviously, the mere fact that some factors are invariant moral reasons would not undermine moral holism, or even ontological holism. But notice that any attempt to offer an example of an invariant reason, moral or non-moral, simply begs the question of how we know other factors could not affect its moral valence—unless, that is, the example is supported by an appropriate argument. Given moral holism, there could be no bare counterexamples—counterexamples not supported by arguments—to the claim that there are no invariant moral reasons. There could only be *arguments* that

---

[11] Hooker claims that "particularism", which he seems to regard as the view that there are *no* invariant moral reasons, is "refuted by counterexamples—that is, by general properties that always count morally in favour of an action, or by general properties that always count morally against [sic]" (Hooker 2000, 7-11). Two of his examples are what we might call simple factors, that φ-ing would benefit others and that φ-ing would promote justice. One might question whether these are invariant moral reasons (as Dancy does), or even whether they are moral reasons at all. His other examples are what we might call complex factors: (1) that φ-ing would produce non-sadistic pleasure, (2) that φ-ing "would involve keeping a morally permissible promise that was elicited from you without coercion or deception", and (3) that φ-ing would harm another outside "the course of justified punishment, or of protecting the innocent, or of respecting their informed wishes, or of pursuing morally legitimate competition" (Hooker 2000, 8, 10, 11). Such complex factors are amenable to holistic analyses, and in each case, it is a substantive question whether the factor in question is a single, complex, invariant moral reason or a complex of factors consisting of a single, variant moral reason and relevant active background factors. Moreover, holistic analyses seem especially attractive for (2) and (3). If he means to claim that these are invariant moral reasons, then Hooker needs arguments to show that they are best understood as single, complex, invariant moral reasons, rather than as complexes consisting of variant moral reasons and relevant active background factors. He also needs arguments to show that their moral valences are indeed invariant.

given factors are such reasons. Moreover, any such argument would have to go beyond an appeal to cases. To think the invariance of a factor's moral valence can be established by an appeal to cases is to miss the chief epistemological import of moral holism: that one cannot learn from some cases—no matter how many—what a factor's moral valence will be in other, dissimilar cases. Thus, what is required of one who maintains a particular factor's moral valence is invariant is an argument that it has the same moral valence in every morally possible case.[12] And given the modality of such a claim, the burden of argument rests with the one who makes it—notwithstanding Shafer-Landau's claim that it is not very difficult to identify some factors that are valence-invariant (Shafer-Landau 1997, 590). Moreover, such an argument would have to show either (a) that the factor in question's moral valence is incapable of varying or (b) that there are no factors capable of affecting its moral valence.[13]

---

[12]Such arguments might, I suppose, take the form of inferences to the best explanation—i.e., that one might argue that it follows from the best explanation of why a factor has the moral valence it does in a range of cases that it has that same moral valence in any morally possible case. But the invariance of the factor's moral valence (i.e., the fact that it is an invariant moral reason) is not a candidate for such an explanation. It is among the *explananda*, not an *explanans*. Cf. Crisp 2000, 36 n. 42.

[13] Unlike Shafer-Landau, Hooker does offer arguments for some of his counterexamples. But what do these arguments show? Some merely purport to show that a given factor's moral valence—e.g., that φ-ing would be stealing's—can be outweighed by another's in some cases. Such arguments clearly offer no support for the claim that the moral valence of the factor in question is invariant. Others purport to show that we should take certain complex facts—e.g., that φ-ing would produce innocent pleasure and that φ-ing would break a morally permissible promise not elicited by coercion or deception—to be invariant moral reasons, rather than certain simple ones—e.g., that φ-ing would produce pleasure and that φ-ing would break a promise. In other words, they purport to show both (a) that certain simple facts are invariantly possessed of neutral moral valence and (b) that it is only more complex facts of which they are parts that have invariant positive or negative moral valence. But these arguments alone show neither. For one thing, to show either that these simple facts are themselves invariantly morally indifferent or that these complex factors are invariant moral reasons, Hooker needs an argument for the claim that there are no morally possible cases in which their moral valences vary. To repeat, there could be

This same objection applies *mutatis mutandis* to Crisp's strategy of appealing to allegedly second-order right-making factors to secure invariance. It simply begs the question of how we know other factors could not affect the moral valence of the (first-order) right-making factors in question. Moreover, given moral holism, there is no *a priori* reason to think that the conditions under which every, or even any, factor's moral valence is invariant are enumerable (Dancy 1993, 81; Dancy 1999b, 26-28; cf. Shafer-Landau 1997, 593-94).[14] Certainly, no thesis about the universalizability of moral judgments or about the supervenience of moral properties that does not simply beg the question against the holist guarantees that they are. And just as appeals to cases cannot establish the invariance of a given factor's moral valence, they cannot establish the conditions under which a factor's moral valence is invariant. And just as anyone who claims a factor's moral valence is invariant owes us an argument, so too does anyone who claims to have enumerated the conditions under which a factor's moral valence is invariant.

---

no bare counterexamples to the claim that there are no invariant moral reasons, only arguments that given moral factors are such reasons. For another, each of the arguments he offers is at least as good an argument for the claim that the simple fact in question is a variant moral reason as it is for the atomistic conclusion Hooker draws. What Hooker needs are arguments that the factors in question are best understood as complex invariant moral reasons, rather than as combinations of variant moral reasons and (some of) the active background factors that can affect their moral valences.

[14] Thus, it is hardly irrational to refuse to engage in an attempt to enumerate such conditions (cf. Crisp 2000, 24-36, 42-44). To the contrary, it is *prima facie* irrational to engage in such an attempt when there is no guarantee it will succeed.

4. Holism and Ontological Atomism

Advocates of ontologically atomistic moral theories may well be unmoved by my

arguments, for they may well believe that the truth of moral holism would be no reason

to think that ontological atomism is false, and hence no reason to think that their

ontologically atomistic moral theories are false. Moreover, one might find a particular

atomistic moral theory plausible for reasons that have nothing to do with atomism of any

sort. For example, one might find act utilitarianism plausible because one finds

teleological accounts of practical reason and impartialist, welfarist theories of value

plausible. And inasmuch as moral holism entails nothing one way or the other about how

probable ontological atomism is, it gives us no reason to believe that act utilitarianism or

any other atomistic moral theory is false.[15] Why, then, should an ontological atomist care

whether moral holism is true or not? After all, if we have confidence in a given atomistic

moral theory, moral holism itself should not undermine that confidence.

This line of thought misses the significance of moral holism for debates within

contemporary moral theory. First, the more plentiful and plausible that the moral holist's

examples are, the more force they have as counterexamples to any atomistic moral theory.

Second, although it is true that moral holism *itself* should not undermine our confidence

in any given atomistic moral theory, it does not follow—and indeed it is false—that moral

holism should not do so. Confidence in any moral theory can only be grounded in

cogent reasoning, and our own reasoning may be tainted in numerous and subtle ways by

---

[15] For simplicity, I assume here that any form of act utilitarianism would be atomistic. One might, I think, question this assumption.

atomistic thinking. For example, we might have accepted atomistic arguments as cogent. If our view about the moral significance of killing verses letting die or the morality of active euthanasia might be informed by Rachels's argument, or indeed if any of our moral views might have been informed by atomistic arguments, then the plausibility of moral holism gives us reason to be less confident in them (cf. Kagan 1988, 28). And it does so even if moral holism itself does not. Likewise, we might have assumed that intrinsic moral relevance is invariable intrinsic moral relevance, or that any evidence against the universality of a statement of the form 'there is a *prima facie* obligation to do A on account of F' (or 'F is a right-making factor') is evidence against its truth. Or we might have relied upon some atomistic version of the thesis that moral judgments are universalizable, or assumed that Newtonian mechanics can serve as a model for moral theory.

Consider, as yet another example, attempts to offer counterexamples to the claim that there are no invariant moral reasons. Typically, these supposed invariant moral reasons are arguably better understood as complexes of right-making and active-background factors than as individual valence-invariant right-making factors—e.g., that ϕ-ing would harm another outside "the course of justified punishment, or of protecting the innocent, or of respecting their informed wishes, or of pursuing morally legitimate competition" (Hooker 2000, 11). Yet, those who offer them invariantly fail to mention, let alone address, this possibility. It is hard not to see such attempts as more evidence of the numerous and subtle ways in which atomistic thinking can taint our reasoning.

Indeed, there really is no way to know whether or how atomistic thinking might have tainted our reasoning. And the real possibility that our reasoning may be tainted by

atomistic thinking in numerous and subtle ways ought to undermine any confidence we

may have in any atomistic moral theory.[16]

Moreover, those of us sympathetic to act utilitarianism and other forms of

consequentialism may have special reason to be concerned that our sympathy rests in part

on reasoning tainted by atomistic thinking. Granted, a consequentialist need not be an

ontological atomist: that the right action is the best action does not entail that the factors

that contribute to the rightness/goodness of actions are valence-invariant ones. And,

granted, atomistic arguments can be used to defend non-consequentialist views (cf.

Kagan 1988, 26-28). But there is a demonstrable history of their being advanced to refute

non-consequentialist views, many of which are also part of commonsense morality. They

have been offered against the do/allow distinction and the killing/letting-die distinction

(Singer 1979, 150-52), against the positive duty/negative duty distinction (Tooley 1972,

30-31), and against absolute prohibitions on certain kinds of conduct (Singer 1979, 195-

97). And they have been used to criticize the view that it is not wrong to save the life of

one person rather than of many, even though it would be wrong not to save one life

_____

[16] It is not clear to me that these remarks would apply in earlier eras. Nor is it clear to me how widely they apply today. Given that Mill was quite aware of the fact that chemical powers do not combine the way mechanical forces do, and of people who thought political economy was more like chemistry than physics in respect of the way the causes it studies interact (Mill 1872, 440; Mill 1873, 169), there would seem to be no reason to suppose that he *assumed* moral powers must combine in just the way that mechanical forces do. (For a relevant and illuminating discussion of Mill's views, see Cartwright 1989, 170-79.) That is, his (supposedly) ontological (and perhaps metaphysical) atomism not withstanding, there would seem to be no reason to presume he was a moral atomist, or that his reasoning was tainted by atomistic thinking. In Ross's case, although he is often taken to be an atomist, there is fairly compelling evidence that he was not just a moral holist but also an *ontological* holist (see Chapter 5). I suspect that much atomistic thinking can be traced to the acceptance of dubious Moorean and Harean views about what is analytic or true given the logic of our moral concepts, views that Mill, Ross, and probably others do not or would not share.

rather than none (Glover 1977, 207-9; Kavka 1979, 291-92).[17] Indeed, if moral atomism

were known to be true, consequentialists could cite the very examples to which moral

holists appeal as evidence against any number of non-consequentialist views on the

grounds that the only satisfactory accounts of why the moral valences of the factors in

question vary as they do hold them to be derivative of (invariant) consequentialist factors

(cf. Pickard-Cambridge 1932; Philips 1987).[18] It is also worth remembering that Moore

and Singer are consequentialists, as is Hare, after a fashion. Thus, if those of us

sympathetic to consequentialism are to be confident in that sympathy, we must—in the

absence of a compelling argument for moral atomism—be confident that it does not rest

on atomistic arguments or other reasoning tainted by atomistic thinking.[19]

---

[17] This view is defended in Anscombe 1967 and Taurek 1977.

[18] Suppose that we did somehow show that moral atomism is true. This alone would not give us any reason to think there are any invariant moral reasons, for it would not give us any reason to think there are any right-making factors. A moral nihilist can be just as fervent a moral atomist as the staunchest moral realist. Moreover, if moral atomism is true, then there are either invariant moral reasons or no moral reasons. Thus, if we did show that moral atomism is true, we would hand the nihilist an argument on a platter: the argument from examples. If moral atomism is true, then the nihilist could cite all the same examples to which moral holists appeal as evidence for moral nihilism, just as consequentialists could cite them as evidence that other moral theories are false. But whereas the consequentialist would claim that the only satisfactory accounts of why the moral valences of the factors in question vary as they do hold them to be derivative of (invariant) consequentialist factors, the moral nihilist would claim that the only satisfactory account of why the *so-called* moral valences of the factors in question vary as they do is that morality is a fiction. Thus, even if we could deprive *the anti-generalist* of her argument by showing that moral atomism is true, we would just hand *the nihilist* an argument for nihilism.

[19] Singer seems to argue that whether a life has intrinsic value depends on whether it is the life of a conscious being (1979, 92). Not only is this position inconsistent with value atomism, but insofar as it implies that the fact that φ-ing would be killing is a variant moral reason, it is also inconsistent with moral atomism. Equally interesting is the fact that Singer does not consider whether the intrinsic value of pleasure might depend in part on whether it is experienced by a self-conscious being rather than a merely conscious one. A holist might regard this as further evidence of how easy it is to fall into the trap of atomistic thinking.

5. <u>Moral Holism and Moral Theory</u>

Finally, I would argue that not only is moral holism weak, plausible, and well-supported by examples, but there are also good theoretical reasons for a moral theorist to be a moral holist.[20] First, moral holism allows for ecumenical conceptions of moral reasons and morally relevant factors. And other things being equal we want ecumenical conceptions of these, ones that are compatible with a wide range of moral theories. Why? Because disputes between such theories are most plausibly understood as disputes within normative moral theory, about what moral reasons or morally relevant factors there are and how they relate to one another. As such, they ought—other things being equal—to be decided by substantive *moral* inquiry, reflection, argument, and theory. We ought therefore to be suspicious of attempts to settle these disputes on grounds that are, from the standpoint of normative moral inquiry, reflection, argument, and theory, *a priori*. And that is just what moral atomism attempts to do. It attempts to exclude moral views and theories positing variant moral reasons from consideration prior to and independently of normative moral inquiry, reflection, argument, and theory (on meta-normative grounds, as it were). It is as if one were to assume without argument that all physical, biological, and other natural events *must* be governed by deterministic laws, and thus that quantum physics must be false, regardless of what our best science tells us—or that the Biblical account of creation is literally true, regardless of what our best science tells us.

_____

[20] The claim here is not that a moral theorist should want to accept that there are variant moral reasons, but rather that there are reasons for a moral theorist to keep the option of accepting such reasons open, if she can.

Second, accounts of at least some substantive views have theoretical advantages

over atomistic alternatives, advantages that any moral theory that offers them will have

over any moral theory that does not or cannot do so.[21] Consider the view that there is a

moral difference between producing sadistic pleasure and producing non-sadistic

pleasure. The challenge for the atomist wishing to hold this view is to show how one can

draw a morally significant distinction between producing the one and producing the other

within an atomistic framework without losing argumentative ground to the moral holist,

for holistic accounts of the moral difference between producing sadistic pleasure and

producing non-sadistic pleasure are more intuitive, more parsimonious, and simpler than

atomistic ones.[22]

To avoid appealing to variant moral reasons in drawing a morally significant

distinction between producing sadistic pleasure and producing non-sadistic pleasure, the

atomist must draw that distinctions in a counterintuitive way. Recall M1 and M2.

Intuitively, it is the fact that McNaughton's nieces and nephews enjoyed the circus that

made it right for him to take them. It is only when we reflect on cases like M2 and our

intuitions about them that we are led to distinguish sadistic from non-sadistic pleasure

and become inclined to say of cases like M1 that it is not the fact that ϕ-ing would

---

[21] In what follows, I assume both that pleasure, at least of the non-sadistic sort, is intrinsically good and that it grounds reasons to promote it. One might, I suppose, deny this, but I take it that would be a high price to pay for rejecting moral holism. Indeed, I take it, many would regard any argument of the form 'view *x* implies non-sadistic pleasure is never intrinsically good and never grounds reasons to promote it' as a *reductio*.

[22] A moral atomist might deny there is any moral difference between these two cases, and that producing sadistic pleasure as such is on a moral par with producing non-sadistic pleasures. But another moral atomist might want to say there is a moral difference between producing sadistic pleasure and producing non-sadistic pleasure (cf. Hooker 2000).

produce pleasure that contributes positively to the rightness of ϕ-ing, but rather the fact that ϕ-ing would or did cause non-sadistic pleasure. Or so I suggest. Now if we accept variant moral reasons we can accommodate both sets of intuitions, for we can distinguish between cases involving sadistic and non-sadistic pleasure in a way that does not require us to abandon either intuition. We can claim the fact that ϕ-ing would produce pleasure is a right-making factor in M1 and similar cases but that it is a wrong-making factor in M2 and cases like it. But if we reject variant moral reasons, then either our intuitions about M2 and cases like it force us to give analyses of M1 and similar cases that are at odds with our intuitions about them or vice versa.[23] Thus, drawing a morally significant distinction between producing sadistic pleasure and producing non-sadistic pleasure without variant moral reasons requires drawing that distinction in a counterintuitive way.

The resulting atomistic view would also be less (ontologically) parsimonious than the one that would result from allowing variant moral reasons. Moral atomists cannot allow that two instances of the same kind can differ in their moral significance. Thus, a moral atomist who distinguishes between sadistic and non-sadistic pleasure and claims that the latter but not the former grounds moral reasons to produce it, or that they ground moral reasons of contrary valence, must draw a distinction in kind between sadistic and non-sadistic pleasure. He must say that sadistic pleasure is a distinct kind of thing from non-sadistic pleasure, a distinct kind of thing with distinct necessary

---

[23] Of course, one has no such problem if one's intuition in cases like M2 is that the sadistic pleasure caused by ϕ-ing *morally* favors ϕ-ing just as any non-sadistic pleasure it caused would. This is not, I take it, a very common intuition, even among those who are willing to accept for theoretical reasons that such an intuition would be correct.

properties that account for the distinct, invariant moral valence of the fact that ϕ-ing would produce sadistic pleasure. In contrast, if a moral holist claims that there is a moral difference between producing sadistic pleasure and producing non-sadistic pleasure, she need posit only one kind of thing, pleasure, which can ground reasons to produce it, reasons not to produce it, or neither, depending on what its contingent properties are (e.g., what its source or object is). The atomist must allow that these other properties exist, but must attribute them to different kinds of things, rather than to different instances of one kind of thing. Thus, the moral holist can draw a morally significant distinction between sadistic and non-sadistic pleasure more parsimoniously than the moral atomist can: she needs to posit one less distinct kind of thing. But this is unsurprising. We should expect that of two alternative accounts of the same phenomena, one that allows that factors can interact with one another to give rise to the phenomena and one that denies this, the former will be more parsimonious than the latter.

The moral atomist who wants to claim that there is a moral difference between producing sadistic pleasure and producing non-sadistic pleasure also faces a dilemma. He must either hold that it is simply a brute fact that there is such a difference or he must offer an explanation of why there is such a difference, an explanation that is distinct from that which the moral holist offers. On the first horn, he surrenders a theoretical advantage to the moral holist, who can appeal to contingent differences between pleasures to explain why there is a moral difference between producing sadistic pleasure and producing non-sadistic pleasure. On the second horn, the moral atomist can avoid losing this advantage only by giving an explanation of the difference that is at least as

attractive as any a holist could offer. But it is not clear the moral atomist can meet this challenge.

The most plausible explanations of why there is a moral difference between producing sadistic pleasure and producing non-sadistic pleasure will appeal to the different sources or objects of these different pleasures, e.g., circuses and the pain of others, and to their different properties. The sources and objects of sadistic pleasures are evil (*viz.*, pain and suffering), while those of non-sadistic pleasures are either neutral or good. And the most plausible explanations of the moral difference between producing sadistic pleasure and producing non-sadistic pleasure will appeal to this fact. But if the atomist offers such an explanation, he bears an additional explanatory burden. He must explain why these differences mark a distinction in kind between sadistic and non-sadistic pleasures. Given that both sadistic and non-sadistic pleasures are pleasures, the atomist owes us an account of why they are nonetheless distinct kinds of things with distinct necessary properties, rather than different instances of the same kind of thing with different contingent properties. In this sense, then, a holistic explanation of the moral difference between sadistic and non-sadistic pleasure is simpler than an atomistic one: it need not account for why the difference between sadistic and non-sadistic pleasure is a difference in kind, because it need not claim that it is such a difference.

Indeed, anytime an atomist draws distinctions, he draws distinctions in kind, and so multiplies the kinds of entities he posits, each with its own set of distinguishing necessary properties, and incurs an obligation to justify doing so and an obligation to give an account of why the distinction is one in kind. Hooker suggest that the atomist can resist the force of the holist's examples by making distinctions. In addition to distinguishing

between producing sadistic pleasure and non-sadistic pleasure, Hooker suggests that we distinguish (a) between keeping promises that are immoral (i.e., inconsistent with one's other obligations) or elicited by coercion or deception[24] and keeping promises that are neither (Hooker 2000, 10) and (b) between doing harm to others "in the course of justified punishment, or of protecting the innocent, or of respecting their informed wishes, or of pursuing morally legitimate competition" and doing harm to others other than in the course of any of these activities (Hooker 2000, 11). That is, he wants to hold that the following are invariant moral reasons: the fact that an act would produce non-sadistic pleasure, "the fact that an act would involve keeping a morally permissible promise that was elicited from you without coercion or deception" (Hooker 2000, 10), and the fact that an act would harm another outside "the course of justified punishment, or of protecting the innocent, or of respecting their informed wishes, or of pursuing morally legitimate competition" (Hooker 2000, 11).

A moral atomist who wants to hold there is a moral difference between keeping the former two sorts of promises must say they are ontologically distinct kinds of things with distinct necessary properties, one of which invariantly grounds a moral reason to keep it and one that does not. Likewise, a moral atomist who wants to hold there is a moral difference between doing the latter two kinds of harm must say they are ontologically distinct kinds of things with distinct necessary properties, one of which invariantly grounds a moral reason not to do it and one that does not. This is unparsimonious,

---

[24] Some may want to say that such promises are promises in name only, that in such cases, no genuine promise is made. But why? Indeed, could it not be that the coercive or deceitful obtaining of a promise is worse than the coercive or deceitful obtaining of benefit for the very reason that it is the obtaining of a genuine promise?

implausible, and demands justification. It also demands an explanation of why the distinction between the two sorts of promises and the two sorts of harm are distinctions in kind, rather than distinctions between members of the same kind that have different contingent properties.

Moreover, this sort of move can quickly lead to absurdity. Suppose that I promise to lend a recently unemployed friend money to pay his rent and that the fact that I did so contributes to the rightness of my lending him the money. Now suppose that my friend subsequently relieves me of the promise because he found a new job. It is still true that I promised to lend the money and that my promise initially constituted or grounded a moral reason for me to do so. But does it still? Surely not. Surely, I now have no more reason to loan him the money than I did before I made the promise. (Indeed, I have less, as he is no longer in need.) Are we to say that unrelieved promises and relieved promises are ontologically distinct kinds of things? Whether my promise was relieved or not, it came into being and was the kind of thing it is when I made it. It couldn't have been a relieved promise when it was made. Did it transubstantiate when it was relieved? No, its prior moral significance was extinguished by the fact that it was explicitly relieved. It gained one contingent property (that of being relieved) and lost another (that of being unrelieved). This is utterly plausible and not in the least mysterious, but a moral atomist must deny that it is even possible. He cannot allow the fact that I promised to lend my friend rent money to be at one time the ground of a reason for me to do so but cease to be such because of a subsequent fact, because he cannot allow that the moral valence of one fact may be affected by another. (Indeed, he cannot allow that such a thing is even conceivable.)

These are all good theoretical reasons to prefer moral holism to moral atomism. Holistic analyses of these views and no doubt many others will be more intuitive, more parsimonious, and simpler than atomistic alternative.[25] And, of course, if a given moral factor is best understood as an invariant one, a holist can allow that it is one. Thus, there are good reasons for a moral theorist to be a moral holist. To all this it may be added that it is quite easy to find examples outside of the moral, or even the normative sphere, of contributing factors behaving holistically. Consider the way the contribution made by a particular spice to the gustatory quality of a culinary dish varies from dish to dish. Cayenne pepper, for example, makes many dishes better, but it would spoil many others. Similarly, whether an additional five knots of wind makes for a better (or a worse) sail depends on many factors, including the current wind speed, the type of boat, the skill of the skipper and her crew, their mood, and the nature of the sail (a competitive race, a leisurely cruise, a fishing trip, etc.). If we accept moral holism, we allow that moral factors could work just the same way, though we allow that they may not. Thus, for those who believe moral factors are factors that contribute to the rightness or wrongness of actions, holism should seem especially plausible.

---

[25] Granted, an atomist might argue that the complexity to which she is committed at this level is compensated for at another, but she would need an argument.

CHAPTER 5

MORAL DISPOSITIONALISM

Moral principles play important roles in diverse areas of moral thought, practice, and theory. We appeal to them in moral deliberation, in justifying our moral judgments and beliefs, and in moral discourse and argument. They play prominent roles in moral learning and education. We regard them as having normative force and, sometimes, try to live by them. They also play a central role in explaining moral phenomena such as the rightness of particular actions, the justness of particular social systems and institutions, and the nature and extent of moral rights. Chapter 1 developed and defended a view—moral generalism—that promises to explain both why moral principles play these roles and why they are capable of doing so. According to moral generalism, moral principles can and do play these roles because true moral principles are statements of general moral fact and general moral facts explain particular moral facts. As we saw in Chapter 3, moral holism maintains that what is a right-making factor, or moral reason to ϕ, in one case may not be one in another, and may even be a wrong-making factor, or a moral reason not to ϕ, given the right context. It holds that the moral valence of a factor is not fixed independently of—and so may be altered by—factors other than itself (active-background factors). Chapter 4 defended moral holism, but we must ask whether these two views, moral holism and moral generalism, are compatible.

Some "moral particularists" maintain that moral holism motivates skepticism about the existence of and need for moral principles, along with skepticism about the viability of principle-based approaches to ethics and moral theory.[1] Here is a sampling of their claims. McNaughton has claimed that "moral principles are at best useless, and at worst a hindrance" to good moral judgment (McNaughton 1988, 190), and that moral holism gives us good reason to give up both the search for general moral principles and the conception of moral theory as such a search (McNaughton 1988, 190-94; McNaughton 2000, 284-86). Dancy claims that moral holism gives us reason to doubt whether there are any moral principles at all, let alone enough to serve as the basis of rational moral thought and judgment (e.g., 1983, esp. 531-32, 533-34; 1993, 66; 2000, 135-36). And Margaret Little claims it gives us "reason to doubt whether there are any moral principles, *even* very complicated ones, capable of codifying the moral landscape" (2000, 277).

I will argue that moral holism is itself a form of moral generalism, one that takes facts about the right- and wrong-making powers of (generic) moral factors to explain certain particular moral facts, namely, the rightness and wrongness of particular actions. I will also argue that a moral-theoretic version of *dispositionalism*—the view that dispositions, powers, or capacities are the fundamental units of explanation—explains

_____

[1] See especially Dancy 1983, esp. 530; Dancy 2004, 78-85; McNaughton 1988, 192-94, 200-201; McNaughton 2000, 284-86; Little 2000. Other moral (or ethical) particularists seem largely motivated by (non-hierarchical) pluralistic theories of the good or the right. They stress that rules and principles can, at best, play a limited role in good moral decision-making. In particular, they stress that moral rules and principles cannot dictate what we should do in particular cases, because numerous, potentially conflicting considerations bear on such decisions, and because no rules are available for sorting out these conflicts. Thus, to make good moral decisions, we require an ability to evaluate particular cases that can neither be reduced to nor represented as rule following. Many of these particularists nevertheless believe that rules and principles can and do play important roles in moral decision-making and other aspects of moral thought, practice, and theory. Ross, for instance, is such a particularist.

both moral holism and moral generalism. In Section I, I argue that certain potential worries about the compatibility of moral holism with moral generalism are unfounded. At the same time, I lay the foundation for showing both that and how moral generalism can readily allow that there are variant moral reasons, and even valence-reversing ones, which I then do in Section II. Section II also lays the foundation for my claim that moral holism is itself a form of moral generalism. Throughout these sections, I use Ross's normative views as illustrative examples, not because I agree with the particulars of his normative moral theory, but rather because discussions of his moral metaphysic and the viability thereof figure prominently in the literature on moral particularism and because—*pace* the conventional wisdom embodied in that literature—he is a moral holist and has the metaphysical resources to combine his moral holism with his moral generalism. Section III, reconstructs (and to some extent reinforces) the reasoning that leads some moral particularists to believe that moral holism motivates skepticism about moral principles and the viability of principle-based approaches to ethics and moral theory. Section IV explains why moral holism is itself a form of moral generalism, why the particularist's argument has no force against either *moral dispositionalism* or a dispositionalist account of moral principles, and why moral dispositionalism explains both moral holism and moral generalism.

1. <u>Moral Holism and Moral Generalism: General Moral Facts</u>

Moral holism presents a *prima facie* challenge for moral generalism. Moral generalism holds that general moral facts (facts about the moral attributes of kinds) explain particular moral facts (facts about the moral attributes of particulars), and it might

seem that if a kind K has a moral attribute M, then any particular K (i.e., any instance of K) must itself have M as an attribute. Granted, the truth of 'Homicide is wrong' does not entail that all homicides are wrong, let alone that any homicide must be wrong (see page 21). But once we distinguish between *prima facie* wrongness and overall wrongness (see Chapter 2), it might seem that this is because what 'Homicide is wrong' says is that homicide is *prima facie* wrong, and one might think that what its being true does not entail is that any homicide must be wrong *overall*. That would be a case of a kind K's having a disposition (*prima facie* wrongness) not entailing that any particular K must manifest that disposition (be wrong overall). But the worry here is that a kind K's having a moral attribute M might entail that any particular K must itself have M—e.g., that homicide's being *prima facie* wrong might entail that any particular homicide must itself be *prima facie* wrong. This is a worry because if it does, and if moral generalism is true, then it might seem as if there could be no variant moral reasons. Indeed, one might wonder whether—if general moral facts explain particular moral facts and any instance of an M kind must itself be M—right-making factors could really have variable moral valences.

Fortunately, this worry is misplaced. The reason why the truth of 'Homicide is wrong' does not entail that all homicides are wrong is not that what 'Homicide is wrong' says is that homicide is *prima facie* wrong. The reason is that 'Homicide is wrong' neither says nor implies that all homicides are wrong. An ascription of an attribute to a kind does not imply that all instances of that kind have that attribute, and hence the truth of such an ascription does not entail that all instances of that kind have that attribute. 'Matches are flammable' (or the synonymous, 'A match is flammable') does not imply that all matches are flammable and so is not falsified by wet matches, which are not. 'Humans have

thirty-two teeth' does not imply that all humans have thirty-two teeth, and so is not falsified by humans who have fewer. Thus, even if what 'Homicide is wrong' says is that homicide is *prima facie* wrong, its being true does not entail that that all homicides are *prima facie* wrong, let alone that any homicide must be so, because 'Homicide is *prima facie* wrong' neither says nor implies that all homicides are *prima facie* wrong.

Two further potential worries remain, however. The first is one about how we are to understand the fact that a kind K has a moral attribute (or any attribute for that matter). How are we to understand the fact that a kind K has an attribute M such that this fact does not entail that all particular Ks have M? The second arises from what some moral generalists (e.g., Ross) hold, in addition to the relatively minimal commitments of moral generalism itself: that actions have the moral attributes they do in virtue of being instances of certain morally significant or relevant kinds. I will address the latter worry first, because doing so will suggest an attractive response to the former one.

Moral theories are typically concerned not just with the question of which kinds and which particulars have which moral attributes but also with the question of why kinds and particulars that have moral attributes have the moral attributes they do. In the case of actions, the latter concern frequently takes the form of a search for the feature of right actions in virtue of which they are right—that is, for the feature of actions that are obligatory overall in virtue of which they are obligatory overall. But it can also take other forms, including the search for the features of *prima facie* obligatory actions in virtue of which they are *prima facie* obligatory. For Ross, *prima facie* obligatory actions are *prima facie* obligatory in virtue of being instances of certain morally significant kinds (Ross 1930, 19). That is, he holds not only that there are *prima facie* obligatory action kinds but also that

particular *prima facie* obligatory actions are *prima facie* obligatory in virtue of being instances of some of these kinds. A moral generalist needn't agree with Ross on this point (see below), but suppose that he is right. Even granting that a kind K's being *prima facie* obligatory doesn't entail that all particular Ks are *prima facie* obligatory, could it be that *prima facie* obligatory Ks are *prima facie* obligatory in virtue of being Ks and yet that some particular Ks are not *prima facie* obligatory?

Yes, it could. And it's a good thing for Ross, too. For despite holding that *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings, he also holds that some promise-keepings are not *prima facie* obligatory. He claims that the *prima facie* obligation to keep a particular promise can be "abolished" or "terminated" by "later developments," namely, the impossibility of keeping the promise and the fact that the promisee no longer wants the promise kept (Ross 1939, 109-10). On this view, circumstances intervening between the time at which one promises to ϕ and the time at which one has promised to ϕ can render one's ϕ-ing no longer *prima facie* obligatory, even though it remains something one promised to do, and hence would still be an instance of promise-keeping. Although Ross doesn't explain how such a thing is possible, given that particular *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings, his appeal to tendencies suggests the answer.

Recall that (for Ross) to be *prima facie* obligatory is to tend to be obligatory overall, or to be disposed to be so. Thus, when Ross claims that particular *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings, he is claiming that they tend (in a dispositional sense) to be obligatory overall in virtue of being

promise-keepings. For a particular promise-keeping to be obligatory overall is (among other things) for it to manifest this tendency. I take it that in addition to this Ross holds the following: that to be *prima facie* obligatory is also to manifest a tendency, namely, a tendency to be *prima facie* obligatory; that membership in certain actions kinds, including the kind promise-keepings, brings this tendency with it; and that when an instance of such a kind manifests this tendency it does so in virtue of being an instance of that kind. This explains how he can hold both that particular *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings and that some particular promise-keepings are not *prima facie* obligatory despite being promise-keepings. Particular *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings because for them to be *prima facie* obligatory is for them to manifest a tendency that being a promise-keeping brings with it—a tendency promise-keepings have in virtue of being the kind of things they are (*viz.*, promise-keepings). And some particular promise-keepings are not *prima facie* obligatory because intervening circumstances sometimes interfere with their manifesting this tendency.

This solves the worry about how a view like Ross's that holds that actions are *prima facie* obligatory in virtue of being instances of certain morally significant or relevant kinds can allow that some instances of those same kinds are not *prima facie* obligatory. It also suggests an attractive response to the worry about how we are to understand the fact that a kind has an attribute such that this fact does not entail that all particular instances of that kind have that attribute. The response is that such facts are facts about the dispositional attributes of kinds—or, more simply, their dispositions. On this account, the fact that homicide is wrong is the fact that homicide is *dispositionally* wrong (i.e., is

disposed, or has a disposition, to be wrong), rather than the fact that homicide is *occurrently* wrong. More generally, general moral facts are facts about the moral dispositions of kinds of actions, states of affairs, beings, institutions, and so on. So understood, the fact that Ks are M does not entail that all Ks are M because it is the fact that Ks are dispositionally M and any particular K will be occurrently M only if and when it manifests the relevant disposition.

This leaves room to disagree over whether moral principles like 'Homicide is wrong' are best understood as saying that homicide and other action kinds are dispositionally wrong (or right) *overall* or as saying that they are dispositionally *prima facie* wrong (or right). That is, there is room to disagree over whether they ascribe tendencies to be right and wrong overall or tendencies to be *prima facie* right and wrong. But for present purposes it doesn't matter which they say or ascribe. Recall that there are two things that could prevent something from manifesting a disposition normal for a thing of its kind—a type-normal disposition (see page 48). An α-condition could prevent it from manifesting that disposition, as the absence of oxygen prevents a match from lighting when struck without affecting its flammability, or a β-condition could prevent it from possessing that disposition, as being wet prevents a match from being flammable (and so from lighting when struck). Given this, even if 'Homicide is wrong' says that homicide is dispositionally wrong overall, its being true would not entail that all particular homicides are dispositionally wrong overall—i.e., *prima facie* wrong. Nor should this be surprising, given what I said about Ross's view. For Ross holds that promise-keepings tend to be obligatory overall yet allows that some particular promise-keepings do not. On his view, promise-keepings tend to be obligatory overall yet certain β-conditions (*viz.*, the one's he

identifies) are capable of preventing particular promise-keepings from manifesting this tendency by preventing them from having it (just as being wet can prevent a match from lighting when struck by preventing it from being flammable).[2]

One might object that it is one thing for homicide's being *prima facie* wrong to not entail that all particular homicides are *prima facie* wrong, but quite another for it to be consistent with some particular homicide's being *prima facie* obligatory. More generally, and less contentiously, one might worry that a kind K's being M entails that particular Ks can't be *contra-M*—i.e., have the attribute that is the contrary of M. But if that were so—if particular homicides couldn't be *prima facie* obligatory because homicides are *prima facie* wrong—how could any homicide be obligatory overall? A particular cannot manifest a disposition that it does not have, and for an action to be obligatory overall is for it to manifest a disposition to be so. A more interesting question is whether it could be that particular *prima facie* wrong homicides are *prima facie* wrong in virtue of being homicides and yet that some homicides are *prima facie* obligatory in virtue thereof.[3] But this is the question of whether being a homicide could be a valence-reversing moral reason, and answering that question is more profitably left to the next section.

_____

[2] Any α- or β-conditions on an action's manifesting a tendency to be *prima facie* obligatory are β-conditions on its manifesting a tendency to be obligatory overall.

[3] More generally, if Ks that are M are M in virtue of being Ks, could there yet be some Ks that are contra-M, rather than M, in virtue of being Ks?

2. <u>Moral Holism and Moral Generalism: Right-making Factors</u>

Thus far, I have been concerned to dispel potential worries about the compatibility of moral holism with moral generalism arising from potential worries about what facts about the moral attributes of kinds entail about the moral attributes of particular instances thereof. In particular, I have tried to dispel the worry that if a kind K has a moral attribute M, then any particular K must itself have M; for if that were so, and if moral generalism is true, then one might wonder whether right-making factors could have variable moral valences. What I want to do now is dispel any residual worries one might have about whether moral generalism and moral holism are compatible by showing that moral generalism can readily allow that there are variant moral reasons, even valence-reversing ones. My argument will also lay the foundation both for a better understanding of why some moral holists are skeptical about moral principles and principle-based approaches to ethics and moral theory and, perhaps paradoxically, for defending my claim that moral holism is itself a form of moral generalism.

The easiest place to begin is with Ross's views about promise-keeping, because—*pace* those who take him to be a moral atomist—they commit Ross to the existence of variant moral reasons. Indeed, the question of whether particular promise-keepings could fail to be *prima facie* obligatory if, as Ross holds, *prima facie* obligatory promise-keepings are *prima facie* obligatory in virtue of being promise-keepings just is the question of whether being an instance of the kind promise-keepings could be a variant moral reason. Thus, because Ross can and does hold that being an instance of promise-keeping makes promise-keepings *prima facie* obligatory but nevertheless fails to make some promise-keepings *prima facie* obligatory, he is *a fortiori* a moral holist. And the β-

conditions that, on his view, prevent some promise-keepings from being obligatory overall by preventing them from being (occurrently) *prima facie* obligatory are, on his view, active-background factors that affect the moral valence of the fact that ɸ-ing is or would be an instance of promise-keeping.[4]

This last—I submit—reveals something important about right-making factors, at least as moral holists conceive them: they are factors with the power to make actions *prima facie* obligatory and, in the case of wrong-making factors, the power to make actions *prima facie* wrong.[5] Recall that moral holists maintain that if a factor *f* fails to have the moral valence in one case that it has in another, then there must be an explanation for this—an explanation in terms of the effects that one or more other, active-background factors have on its moral valence. One might wonder whether a moral holist can insist on this. She can, if she holds that a right-making factor *f* disposes actions of a corresponding kind K to be *prima facie* obligatory and must manifest this disposition in suitable circumstances unless something—some relevant α- or β- condition—prevents it from doing so.[6] On this view, having a tendency (to use Ross's term) to make actions of a kind K *prima facie* obligatory, or to *prima facie* obligate agents to perform actions of that

---

[4] Ross considers Pickard-Cambridge's claim that there are numerous other such factors and finds them wanting, but in each case Ross argues that Pickard-Cambridge is mistaken in thinking that the factor he identifies is such a condition (Ross 1939, 94-98; Pickard-Cambridge 1932, 158, 163-66). At no point in his argument does Ross suggest that Pickard-Cambridge is mistaken in thinking that there could be such conditions. (Mostly, Ross thinks the promises in question are implicitly conditional—promises to do *x* if [or unless] *y* obtains—and that a relevant condition fails to [or does] obtain.)

[5] It might be more accurate to say that they are such powers, and what I say here is meant to be consistent with the view that they are.

[6] The qualification 'in suitable circumstances' is needed here because some dispositions (e.g., fragility) manifest only when "triggered".

kind, is *constitutive* of being a right-making factor, and having a tendency to make actions of a kind K *prima facie* wrong, or to *prima facie* obligate agents to not perform actions of that kind, is constitutive of being a wrong-making factor. This is what it is to be a right-making factor and a wrong-making factor, respectively; and it explains why if a factor *f* fails to have the moral valence in one case that it has in another there must be an explanation for this in terms of the effects that one or more other, active-background factors have on its moral valence.

This dispositional account of right-making factors also illuminates the moral holist's claim that what is a right-making factor or moral reason to ϕ in one case may not be one in another. Just as we must distinguish between an action's having a tendency to be obligatory and its manifesting that tendency, we must distinguish between a factor's having the *power* to make actions obligatory and its *exercising* that power in particular cases. For a factor to have a positive deontic valence in a case is for it to exercise a right-making power in that case, for it to have a negative deontic valence in a case is for it to exercise a wrong-making power in that case, and for it to have no deontic valence in a case is for it to exercise neither a right- nor a wrong-making power in that case. On Ross's view, for example, being an instance of promise-keeping is a right-making factor—it has the power to make particular promise-keepings *prima facie* obligatory, which consists in its tendency to make instances of promise-keeping *prima facie* obligatory. But it does not always exercise this power, because sometimes other factors prevent it from doing so. The moral holist's claim that what is a right-making factor in one case may not be one in another is the claim that a factor that exercises or manifests a right-making power or tendency in one case may fail to do so in another.

What of the seemingly more radical claim that at least some moral holists make: that what is a right-making factor in one case may be a wrong-making factor (and vice versa) given suitable circumstances? We can now see that this is simply the claim that factors can have multiple dispositions that manifest in different circumstances. Thus, the claim that the fact that φ-ing would produce pleasure is a valence-reversing moral reason is simply the claim that it has both the power to make φ-ing *prima facie* obligatory and the power to make φ-ing *prima facie* wrong, and that which of these powers it exercises depends on other factors, such as whether the pleasure is an innocent pleasure or a sadistic one. And I can see no reason why even Ross—who holds that actions are *prima facie* obligatory and *prima facie* wrong in virtue of being instances of morally significant kinds—cannot allow that there are such reasons. The same molecular structure that disposes ripe lemons to look yellow in white light disposes them to look black in blue light. The same molecular structure that disposes water to nourish me disposes it to kill me if I drink or inhale too much of it. The same adaptation that disposes survival in one environment disposes extinction in another. I see no reason why being a promise-keeping could not carry with it both a disposition to be *prima facie* obligatory and a disposition to be *prima facie* wrong. Nor do I see any reason why the latter could not manifest, rather than the former, in certain kinds of cases—e.g., those in which the relevant promise was one to do something nefarious. Thus, I see no reason why being a promise-keeping or a pleasure-producing action could not be a valence-reversing moral reason. Indeed, I see no reason why a moral generalist could not hold that *all* moral reasons are such reasons, including—to answer the question raised at the end of the

previous section—being a homicide. Hence, I see no reason why a moral generalist must

hold that there are, or need to be, any invariant moral reasons.

3. Moral Holism and Moral Particularism

The claim that a moral generalist could hold that all moral reasons are valence-

reversing reasons may be thought surprising. Chapter 1 argued that moral generalism

promises to explain why moral principles can and do play the important roles that they

play in diverse areas of moral thought, practice, and theory. But some "moral

particularists" maintain that moral holism motivates skepticism about the existence of

and need for moral principles, along with skepticism about the viability of principle-based

approaches to ethics and moral theory. Why should the truth of moral holism lead us to

be skeptical of moral principles and principle-based approaches to ethics and moral

theory? Here I will focus on two specific theses that some particularists suggest moral

holism motivates. The first is that it is astronomically unlikely that morality is principled,

or that it would be "a cosmic accident" if it were so (Dancy 2004, 82; cf. McNaughton

2000, 194; Little 2000, 277). The second is that, although there are moral reasons or

right-making factors, there are no moral principles (Dancy 1983, 530; McNaughton 2000,

194; Little 2000, 277).

I take it that the particularist's argument is this:

1. *Moral holism*: what is a right-making factor in one case may be no such
   factor or (even) a wrong-making factor in another case; it may have a
   different—even the contrary—moral valence in another case, because
   its moral valence may be affected by factors other than itself.

2. *The entailment thesis*: any moral principle p entails that a corresponding
   right-making factor f has invariable moral valence.

3. Therefore, morality is not principled necessarily: right-making factors do not entail corresponding moral principles.

4. Moreover, morality is not principled in fact: it is not the case that for each right-making factor *f* there is a corresponding moral principle *p*.

5. Moreover, there are no moral principles.

This argument is much stronger than it appears—stronger even than some particularists appear to think. But though it posses a considerable challenge to those who think that morality is codifiable in the ways that some suppose it to be, it undermines moral generalism not at all.

Much of the force of this argument lies in the challenge it puts both to those who assume that morality is principled and to those who assume that moral principles are important elements of moral thought, practice, and theory. (3) follows from (1) and (2): if moral holism and the entailment thesis are both true, then morality is not principled necessarily. Indeed, if they are both true, then there need not be *any* moral principles. Nevertheless, most moral theorists assume that one aim, if not the aim, of moral theory is a principled moral theory. If morality is not principled necessarily, what justifies this assumption? Certainly not the existence of a few seemingly plausible moral principles. And if there need not be any moral principles, what justifies the assumption that moral principles are important elements of moral thought, practice, and theory? Arguments are owed here, arguments in defense of the assumptions that underlie principle-based approaches to ethics and moral theory. But establishing this does not exhaust the force of the particularist's argument.

Neither (4) nor (5) strictly follows from (1) and (2). But it is *prima facie* irrational for anyone who accepts both moral holism and the entailment thesis to deny either (4) or

(5)—i.e., to believe that morality is, in fact, principled or even just that there are some moral principles. If the entailment thesis is true, morality is principled only if there are *only* invariant moral reasons. But it is *prima facie* irrational for a moral holist to think that this is even likely to be the case, let alone that it is so (cf. Dancy 1993, 66; Kagan 1988, 28). If moral holism is true, what reason do we have to think it true, or even likely, that there are *no* variant moral reasons? If moral reasons as such need not be valence-invariant, then what reason is there to believe that there are no such reasons, or that all moral reasons can be reduced to valence-invariant ones? For that matter, what reason is there to believe that there is even one whose moral valence is invariant? Absent a satisfactory answer to the former question, it is irrational for a moral holist to believe there are only invariant moral reasons, and thus for anyone who accepts both moral holism and the entailment thesis to believe that morality is principled. Absent a satisfactory answer to the latter question, it is irrational for anyone who accepts both theses to believe either that there are any moral principles or, *a fortiori*, that principle-based approaches to ethics and moral theory are justified.

Once again, the force of this argument is not exhausted by its showing that certain assumptions are in need of defense. If right-making factors are indeed factors that have a tendency to make actions of a kind K *prima facie* obligatory, then (*a*) we ought to be skeptical about the existence of invariant moral reasons, (*b*) the burden of argument with respect to any claim that the moral valence of a factor is invariant rests with the one who makes it, and (*c*) that burden is a heavy one to bear. For any given disposition, it is a substantive question whether there are conditions that can prevent it from manifesting (Mumford 1998, 86-86). Moreover, a great many dispositions are subject to such

conditions, and it may be that all are. Thus, good methodology treats all dispositions as subject to such conditions, as is arguably witnessed to by the care that scientists take to control for interfering factors when designing experiments and interpreting observational and experimental results. The moral holist simply claims that right-making powers, wrong-making powers, and other moral dispositions are no exception, and I know of no argument, compelling or otherwise, that purports to show that she is wrong to do so. Thus, we ought to be skeptical about the existence of invariant moral reasons, and the burden of argument with respect to any claim that the moral valence of a factor is invariable rests with the one who makes it, not with the one who is skeptical of it.[7] Such skepticism seems all the more reasonable, and the invariabilist's burden all the greater, given the extent to which *we* can and do interfere with the manifestation of various dispositions. Doctors and engineers interfere with the manifestations of biological and physical dispositions routinely. And I take it that one of the things that Dancy's Contraband example purports to show is that we can interfere with the manifestations of moral dispositions, too. Thus, showing that a factor's moral valence is invariable requires showing not only that its tendency to make actions *prima facie* obligatory could not be interfered with by anything, but also that it could not be interfered with by anyone.

---

[7] It is largely for this reason that I don't find either McNaughton and Rawling's claim that "when it comes to certain thick [moral] concepts, the onus is on the thin intuitionist to justify her claim that the valence *can* vary" (McNaughton and Rawling 2000, 273) or Little's much stronger claim that "thick moral features . . . are guaranteed of carrying a given valence of moral significance" (Little 2000, 289) at all convincing. Nor is it obvious that being just, cruel, and so one, are really right-making factors, rather than, say, ways of being right or wrong, or otherwise morally or ethically reasonable or unreasonable.

Moreover, even if this could be shown for one or some moral reasons, that would not be nearly enough to justify principle-based approaches to ethics and moral theory.

By far, then, the most promising response to the particularist's argument is to deny the entailment thesis. One way to do this is to claim that we can simply append exception clauses to our (initially atomistic) moral principles as needed, such that instead of saying that certain factors always or invariably have certain moral valences, our principles say that certain factors always or invariably have certain moral valences under certain enumerated conditions.[8] Quasi-atomistic variants of this response claim that we should include any relevant active-background factors in a 'complete' or 'full' specification of a moral reason (Hooker 2000, 14) and that we should speak of complexes of right-making factors and (other) factors relevant to their weights and valences as "having *prima facie* weight *as wholes*, and of their constituents as individually having certain weights *in the absence of* the other constituents" (Crisp 2000, 47 n. 73). Although such responses might look attractive, at least initially, because they do not make the existence of moral principles contingent on the existence of invariant moral reasons, at least not simple ones, they ultimately fail.

One problem is that such responses make the existence of moral principles contingent on the enumerability of the conditions under which the moral valences of factors are invariable. As I argued in Chapter 4, there is no *a priori* reason to think that the conditions under which every, or even any, factor's moral valence is invariable are

---

[8] Cf. Audi 1996, 119-20; Hooker 2000, 7-8, 11; Jackson, Pettit, and Smith 2000, 96-99; McNaughton and Rawling 2000, 266-72; Shafer-Landau 1997, 593-94.

enumerable—at least not given moral holism.[9] And just as no non-question-begging

thesis about the universalizability of moral judgments or the supervenience of moral

properties guarantees that they are enumerable, no such thesis guarantees that there are

any principles of a sort capable of playing any of the roles that moral principles play in

moral thought, practice, and theory, let alone that morality is principled in any meaningful

sense. Nor is such an enumeration guaranteed to be nearly so simple a task as some

opponents of particularism suggest. The latter often suggest that one can capture claims

about, say, the variable valence of the fact that ϕ-ing would produce pleasure in principles

such as 'The fact that ϕ-ing would produce pleasure is always a reason to ϕ unless the

pleasure would be a sadistic one, in which case it is never a moral reason to ϕ (or always a

moral reason not to ϕ)'. But what holism maintains is not that the moral valence of a

factor may *always* be affected by other factors *whenever* they are present, but rather that the

moral valence of a factor *may* be affected by other factors. Whether or not one factor

affects the moral valence of another may itself depend on what other factors are present

along with it. Thus, that a factor has the ability to affect the moral valence of another

factor does not guarantee that it will do so whenever it is present, nor even that it will

usually do so when it is present. Thus, the fact that a pleasure's being sadistic can alter

the moral valence of the fact that ϕ-ing would produce pleasure does not guarantee that it

will usually do so, let alone that it will always do so.[10] Moreover, to enumerate the

---

[9] Cf. Dancy 2004, 81-82; Kagan 1998, 184-85; Shafer-Landau 1997, 594.

[10] For similar reasons, we cannot assume that how much valence-variability there is depends on how right-making factors are, or should be, differentiated such that the finer such factors are differentiated (i.e., the more specific they are), the less variability there is. For example, we cannot assume that there would be less variability if the relevant factors are that ϕ-ing would produce

conditions under which a factor's moral valence is invariable would be to enumerate the conditions under which it manifests a particular disposition (e.g., a right-making tendency), and any such enumeration would, at best, be subject to revision based on new information. Although moral particularists are often thought to hold that the facts about which actions bear which moral properties are too complicated to be codified, the thrust of this argument is different. It is, rather, that the facts about how actions come to have the moral properties they do are such that—at best—it seems highly unlikely that the facts about which actions bear which moral properties could ever be definitively written down, even in part.

Let me reiterate what I take to be at stake here. Jackson, Petit, and Smith—who describe themselves as "principle-ists" (2000, 99)—appear content to argue that there must be "a descriptive pattern uniting [all] the right acts" (89), a pattern that "must be codifiable *in principle* [my emphasis]," whether or not it is ever "codified, in the sense of being 'before the mind', when a moral judgement is made" (91).[11] And Kagan maintains

---

sadistic pleasure and that φ-ing would produce non-sadistic pleasure, rather than that φ-ing would produce pleasure. There are cases in which it is plausible to say the fact that φ-ing would give someone pleasure whose object is the suffering of another is morally irrelevant. Here is one. Stan is hiking with Kenny and Eric. Kenny is accident prone, and Eric is prone to *Schadenfreude*. Kenny breaks his leg. Setting it will cause Kenny pain, which will give Eric *Schadenfreude* (even more than he is already experiencing). Plausibly, this last fact is simply irrelevant to what Stan ought to do. It is also arguable that the fact that φ-ing would produce sadistic pleasure can be a moral reason to φ in certain contexts—e.g., in the context of *certain* sadomasochistic sexual relationships. How to differentiate right-making factors (and other morally significant kinds) is an important and complex methodological question, which is beyond the scope of this work. But even if the extent of valence-variability depends on how finely or coarsely right-making factors are differentiated, there is no reason to assume that finer differentiation entails less variability—or even less reversibility. In fact, the more finely differentiated right-making factors are, the more numerous are the factors whose moral valences may vary.

[11] Dancy 1999a is a response to their argument.

that "universalists" need only contend that "a complete set of interaction principles" exists that "governs all possible interactions of normative factors," in the sense that its members "correctly assign a determinate moral status to the relevant actions on the basis of the factors at play" (Kagan 1998, 186). "They need not," he emphasizes, "believe that we can discover these principles or articulate them" (186). Moreover, he insists, "they certainly need not believe that exceptionless interaction principles will have any role in practical deliberation" (186). Perhaps some who think of themselves as moral generalists will be content with such principle-ism or universalism. But I want to defend the claim that moral principles can and do play important roles in diverse areas of moral thought, practice, and theory. To do that, one must offer an account of moral principles on which they are capable of playing those roles—especially if one appeals, as I do, to the fact that moral principles play such roles in order to defend one's interest in, the need for, or the existence of such principles and principle-based approaches to ethics and moral theory. And it is not unreasonable to think that there is a limit to how complicated and qualified moral principles could be and still play the roles they play. Nor is it unreasonable to think that they must be discoverable in order to do so.

But even putting aside such concerns about codifiability, there is a second, independent, and more fundamental problem with these and other extant responses to the particularist's argument. They rest on a familiar, but mistaken assumption: that moral principles are (or would be) statements of *deontic laws*, regular associations of right-making factors and either *prima facie* rightness or overall rightness. Many assume that moral principles are or would be such statements, statements such as 'An action is right if and only if it maximizes utility', or 'All promise keepings are *prima facie* right', or "That an

act would involve keeping a morally permissible promise that was elicited from you without coercion or deception always counts morally in favour of your doing the act" (Hooker 2000, 10). But this is a mistake. Qualified or not, neither statements of deontic laws nor such laws themselves could be moral principles, because they lack the explanatory power of moral principles.

Moral principles purport to explain why particulars bear the moral properties they do, and true principles of obligation (or the facts they state or express) explain why particular obligatory actions are obligatory.[12] 'One ought to keep one's promises' (or the fact that one ought to do so) explains why I ought to meet Andy for a pint, given that I promised to do so. But even if 'All promise keepings are obligatory' were true, it would not explain why this or any particular promise-keeping is obligatory. It is incapable of doing so, just as 'All men are mortal' (or the fact that all men are) is incapable of explaining why Socrates, or any other man, is mortal. The problem is that statements of deontic laws are universal statements and—as I argued in Chapter 1—the truth of a universal statement—although it implies the truth of its instances—is utterly irrelevant to *why* its instances are true. That two classes (e.g., the class of right acts and the class of maximally utile acts) are coextensive is utterly irrelevant to why any particular member of one is a member of the other (e.g., to why any particular maximally utile act is right). That all the members of a certain class (e.g., the class of promise-keepings) possess an attribute (e.g., obligatoriness) is utterly irrelevant to why any particular member of that

_____

[12] Indeed, this is one of the principle roles that moral principles play in moral thought, practice, and theory, and it is why they can and do play many of the other roles they play therein (Chapter 1).

class (e.g., my meeting Andy) possesses that attribute. And, we can now add, that a given factor (e.g., that φ-ing would involve keeping a morally permissible promise that was elicited without coercion or deception) always counts morally in favor of φ-ing, or always does so under certain conditions, is utterly irrelevant to why it counts in favor of any particular φ-ing, and hence to the moral status of any φ-ing that it counts in favor of. Regardless of their modality, deontic laws would not be relevant to why particulars have the moral properties they do. To see this, one need only imagine asking someone who believes that it would be wrong for two women to marry for an explanation of why this would be wrong, and being told that (necessarily) all same-sex marriages are wrong, or that (necessarily) the fact that a marriage would be one between members of the same sex always counts morally against it. Only someone who did not understand the question would offer such an answer. Perhaps conceiving of moral principles as universal statements of one kind or another is appropriate for non-cognitivists and other anti-realists, but if moral principles explain why particulars bear the moral properties they do, then they are neither universal statements nor facts expressed by such statements. Hence, they are neither deontic laws nor statements thereof.

4. Moral Holism, Moral Generalism, and Moral Dispositionalism

In Chapter 1, I argued that (unlike universal moral statements) statements of general moral fact have the explanatory power of moral principles, as well as the modal force thereof. We can now see why this is so. A statement of general moral fact, which ascribes a moral attribute to a kind, implies that particulars of that kind must and will have that attribute in the absence of interfering factors—factors that interfere with their

having that attribute, either by preventing them from having that attribute or by causing

them to have a different, incompatible attribute (e.g., to be occurrently right overall,

rather than occurrently wrong overall). We can now see that this is because general moral

facts are facts about the moral dispositions—including the moral tendencies and

powers—of kinds of actions, states of affairs, beings, institutions, and so on.[13] Likewise,

it is because general moral facts are dispositional facts that statements of such facts can

explain particular moral facts. That homicide is wrong can explain why *that* homicide is

wrong because for homicide to be wrong is for it to be dispositionally wrong and for a

particular homicide to be occurrently wrong is for it to manifest a disposition to be

wrong. Likewise, the principle 'Homicide is wrong' can explain why *that* homicide is

wrong because what it says is that homicide is dispositionally wrong (i.e., because it

ascribes a disposition or tendency to be wrong to homicide), and because to be

occurrently wrong is to manifest a disposition to be wrong.[14]

All this is perfectly consistent with moral holism. In fact, moral holism is itself a

form of moral generalism. As I defined it, moral generalism is the view that general

moral facts explain particular moral facts. Moral holism claims that right-making factors

explain why particular right actions are right, and that wrong-making factors explain why

particular wrong actions are wrong. And it conceives of these as generic factors (states of

affairs, features of actions, etc.) with the power to make actions right and wrong. In

other words, it conceives of them as *kinds* of factors with dispositional attributes that

---

[13] Thus, interfering factors are α- and β-conditions.

[14] It is also for this reason that (unlike moral imperatives) general moral facts are sufficiently
fundamental to be moral principles.

explain certain particular moral facts, namely, the rightness and wrongness of particular actions. And facts about the right-making powers of kinds of factors are themselves general moral facts, facts about the moral attributes of kinds. Thus, moral holism is itself a form of moral generalism.

Notice that a moral holist need not deny either that certain kinds of actions are (dispositionally) right or that a particular K's (occurrent) rightness can be explained by the (dispositional) rightness of the kind K. One might hold that certain kinds of actions tend to be right because certain (kinds of) factors are right-making ones. On this view, it would not be the rightness of the kind promise-keepings that explains why, say, the fact that one promise to ϕ is a right-making factor. Rather, the latter would explain the former. Yet, one could still cite the obligatoriness of promise-keeping to explain why a particular obligatory promise-keeping is obligatory.

Notice, too, that a kind is not *prima facie* right or wrong simply because its instances are capable of being so. Suppose the fact that ϕ-ing would be lying is a valence-reversing moral reason—that it has both the power to make lies *prima facie* wrong and the power to make them *prima facie* obligatory (as Dancy claims it does in the game of Contraband). It does not follow that lying is *prima facie* obligatory. In fact, that type-normal lies are *prima facie* wrong is inconsistent with the *prima facie* obligatoriness of lies *qua* kind. For that lying is *prima facie* obligatory would entail that type-normal lies are *prima facie* obligatory. Thus, to answer the other question raised at the end of Chapter 2, neither the dispositional account of *prima facie* obligations nor moral generalism requires that every particular *prima facie* obligatory action be an instance of a *prima facie* obligatory action kind.

What I am sketching here is a moral-theoretic version of *dispositionalism*, the view

that dispositions, powers, or capacities—rather than laws—are the fundamental units of

explanation (see, e.g., Mumford 1998, chap. 10; Cartwright 1999, esp. chaps. 3 & 4).

According to *moral dispositionalism*, moral dispositions, powers, or tendencies—rather than

laws—are the fundamental units of *moral* explanation,[15] which implies that particular

moral facts—facts about the moral properties of particulars—are manifestations of such

dispositions.[16] (If one takes the 'laws' of a domain to be the fundamental units of

explanation therein, whatever those might be, then one may take moral dispositionalism

to be the view that moral laws are dispositions—rather than, say, regular associations of

events or properties.)

Now the kind of moral particularism that is skeptical about moral principles—and

principle-based approaches to ethics and moral theory—presupposes moral

dispositionalism, for it holds that right-making factors conceived as factors with the

power to make actions *prima facie* right and *prima facie* wrong are the fundamental units of

moral explanation. But this just shows that the particularist's skepticism about moral

principles *as such* is misplaced. For moral dispositionalists need not, and should not, be

skeptical about moral principles *properly understood*. To be sure, at least some

dispositionalists are skeptical about the existence of *laws* conceived as regular associations

---

[15] So stated, moral dispositionalism does not preclude explanations of moral dispositions, because such explanations would not themselves be *moral* explanations. Rather, they would be explanations of morality itself. Thus, a moral dispositionalist could be an ethical naturalist, non-naturalist, rationalist, or what have you.

[16] Moral dispositionalism in this sense is not the view that moral properties are response-dependent properties (dispositions to produce or elicit responses in or from agents or valuers); nor does it entail that they are such properties.

of events or (occurrent) properties (e.g., Cartwright 1999, 49-50). But dispositionalists

need not, and frequently are not, skeptical about the *principles* that are so commonly

regarded as laws so conceived—including the principles of well-confirmed scientific

theories. Rather, they offer an alternative account of these principles—an account they

argue is truer to the roles that principles play in scientific thought, practice, and theory.

Dispositionalists hold that natural "laws" should not be understood as regular

associations of events or occurrent properties, but rather as descriptions of capacities or

ascriptions of dispositions (e.g., Cartwright 1999; Lowe 1989, chap. 8; Mumford 1998,

chap. 10; Mumford 2000). Cartwright, for instance, argues that Newton's law of gravity

($F=Gm_1m_2/r^2$) describes the capacities of massive bodies *qua* massive and that Coulomb's

law of electrostatic attraction and repulsion ($F=q_1q_2/4\pi\varepsilon_0r^2$) describes the capacities of

charged bodies *qua* charged—or what it is in their natures to do *qua* massive and *qua*

charged, rather than their actual behavior (1999, 65, 82).[17]

Dancy charges:

> Moral principles, however we conceive of them, seem all to be in the
> business of specifying features as *general* [i.e., invariant] reasons. The
> principle that it is wrong to lie, for instance, presumably claims that
> mendacity is always a wrong-making feature wherever it occurs. (Dancy
> 2000, 135; Dancy 2004, 76)[18]

---

[17] Compare this with Ross on gravitation and the laws of nature: "*Qua* subject to the force of gravitation towards some other body, each body tends to move in a particular direction with a particular velocity; but its actual movement depends on *all* the forces to which it is subject. It is only by recognizing this distinction that we can preserve the absoluteness of laws of nature" (1930, 28-29). It is not clear whether Ross thinks it necessary that laws of nature be absolute, nor which (if either) he thinks fundamental, laws or tendencies.

[18] Cf. Dancy 1993, 66; Little 2000; McNaughton and Rawling 2000.

But a moral dispositionalist can offer a conception of moral principles on which they have *nothing* to do with invariant moral reasons. And she can do this because she can conceive of moral principles as statements of general moral fact, and because she can conceive of such facts as dispositional ones.

What the moral dispositionalist can say—and what the moral holist should say—is that moral principles ascribe moral tendencies or other moral dispositions to sorts or kinds (of actions, beings, and so on).[19] That is, she can say that the attributes that moral principles ascribe to kinds are dispositions. Thus, she can say that 'Lies are wrong' should be understood as 'Lies are *dispositionally* wrong' rather than as 'Lies are *occurrently* wrong', that 'One ought to keep one's promises' should be understood as 'Agents are *dispositionally* obligated to keep their promises' rather than as 'Agents are *occurrently* obligated to keep their promises', and so on. So understood, such principles ascribe *dispositions* to sorts or kinds; but individuals thereof are *occurrently* obligatory, wrong, or obligated only if and when they possess and manifest these (type-normal) dispositions. So understood, moral principles are not in the business of specifying features as invariant moral reasons. Nor do such principles imply that corresponding right-making factors cannot exercise contrary moral powers, or have contrary moral effects, in suitable circumstances. Thus, for example, 'Lying is wrong' does not imply that the fact that ϕ-ing would be lying cannot make ϕ-ing *prima facie* right, rather than *prima facie* wrong, in suitable circumstances (e.g., in a game of Contraband).

---

[19] Cf. Lowe 1989, chap. 8; Mumford 2000, 271-72, 278-79. Lowe 1989, chap. 9 offers a formal sortal logic.

There is, to be sure, a difference between such paradigmatic moral principles and paradigmatic scientific "laws". Paradigmatic moral principles are more like 'Ravens are black' or 'Masses attract each other' than they are like '$F = Gm_1m_2/r^2$'. But this is to be expected, as the former are principles of commonsense moral thought rather than moral theory, and it is far from a problem for moral dispositionalism. On a dispositionalist construal, 'Ravens are black' and 'Masses attract each other' ascribe dispositions to kinds, while '$F = Gm_1m_2/r^2$' describes and quantifies the capacity of massive bodies *qua* massive to attract one another. Moral theories may well include principles similar to the latter that quantify the moral tendencies of, say, promise-keepings *qua* promise-keepings or the moral powers of, say, acts of promising. Indeed, Mill's "Actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness" (Mill 1871, 2.2) might be read as just such a principle (as '$R \propto H$', assuming that H can take negative values).[20] Short of such quantification, moral-theoretic principles might describe moral tendencies in greater detail than do the paradigmatic moral principles of commonsense morality. Because formulating either sort of principle would require specific knowledge of morally relevant factors, their moral powers, and the ways in which they can and do interact, a moral dispositionalist is likely to think that the

---

[20] Thanks to Evan Tiffany for this formulation of Mill's principle. A dispositional interpretation of this particular claim may very well not be the best one; but Cartwright's dispositionalism is inspired in no small part by her reading of Mill on tendencies, and it is an interesting question whether Mill ought to be read as a moral dispositionalist. Although the version of moral dispositionalism that I sketch here builds on Ross's view in a number of ways, moral dispositionalism itself is an abstract metaphysical view with little or no normative content. Thus, moral dispositionalism is not committed to Rossian pluralism. Indeed, Ross himself seems to have understood utilitarianism in a dispositionalist way, as holding that promise-keepings and all other actions that tend to be obligatory do so *qua* utile (Ross 1939, 82).

chief task of normative moral theory is to attain such knowledge: to identify (and

differentiate) morally relevant factors, determine their moral powers, and discover how

they can and do interact.

None of this, it should be said, entails that moral dispositionalists must deny that

there are deontic laws, or that morality is codifiable in terms of such laws. But they will

maintain that if there are any such laws, and if morality is so codifiable, it is only as *a*

consequence of the ways in which morally relevant factors can and do interact with one

another in particular cases. In other words, morally relevant factors are responsible for

whatever deontic laws or codifiability there might be, rather than *vice versa*.[21] Short of that,

a moral dispositionalist might maintain that there are lots of deontic laws, but insist that

they hold only *ceteris paribus* or in ideal conditions—that is, in conditions suitable for the

manifestation of the relevant moral dispositions. But whatever her stance on deontic

laws, or on moral holism's implications for the existence of such laws or the codifiablity

---

[21] Kagan expresses just the opposite view: "If the moral status of any given act is indeed determined by the various normative factors at play, then it seems that there must be some principles or laws—governing the interactions of those factors—by virtue of which the particular combination of factors results in the particular moral outcome that it does" (Kagan 1998, 184). Given moral holism, and given that we may be incapable of knowing or even articulating the laws Kagan claims must exist (both of which Kagan grants), one wants to know, first, why there must be such laws and, second, why they must govern, rather than be governed by, the normative factors at play. Kagan's reason for thinking there must be such laws appears to be that "the moral status of any given act is indeed determined by the various normative factors at play." But, as he seems to allow, this doesn't guarantee anything but case-specific laws—laws "governing" cases "*exactly* alike in every single aspect" (184), and one wants to know what statements of such laws would say, other than that normative factors interact in particular cases in the way that they interact in particular cases. Moreover, if it really is *the various normative factors at play* that determine the moral status of any given act, what work is there for Kagan's laws to do? They might determine how the factors interact; but why must anything do that? (Keep in mind, that we could just as well ask what determines the laws that determine how the factors interact.) And even if something must determine how the factors interact, why must it be anything other than the factors themselves and the powers they are or possess? Finally, Kagan's view suffers from all the defects of the enumeration strategy I discussed at the end of Section 3.

of morality in terms thereof, a moral dispositionalist neither need nor should conceive of moral principles as such laws.[22]

Nor should a moral dispositionalist regard moral holism as a problem for the view that moral principles are capable of playing the important roles that they play in moral thought, practice, and theory. The particularist's argument—resting as it does on moral holism and the entailment thesis—has no force against either moral dispositionalism or a dispositionalist account of moral principles. First, far from being inconsistent with moral holism, moral dispositionalism explains moral holism: the moral valence of a factor may be affected by factors other than itself because for a factor to have a moral valence in a case is for it to exercise a right- or wrong-making power in that case. Second, on a dispositionalist account of moral principles, the entailment thesis is simply false, and the conditions under which the moral valences of factors are invariant need not be enumerable in order for there to be moral principles. Moreover, moral dispositionalism explains why moral principles can and do play the roles that they play in moral thought, practice, and theory. For moral dispositionalism explains moral generalism: general moral facts explain particular moral facts because the former are facts about the moral

---

[22] Thus, while I agree with Pietroski that moral principles play the same explanatory role in normative moral theory that what he and (some) others take to be *ceteris-paribus* laws play in the special sciences (cf. Pietroski 1993, 499), I reject his account of moral principles, according to which moral principles are *ceteris-paribus* laws of (roughly) obligation overall (Pietroski 1993, 400, 499; cf. Brink 1994b, 216-17, 245). I would also reject the suggestion that moral principles are *ceteris-paribus* laws of *prima facie* obligation, as well as Pietroski's suggestion "that force talk is interchangeable with generalization talk" (Pietroski 1993, 499), since I take forces and other dispositions to explain, and to be responsible for, whatever generalizations might hold, be they *ceteris-paribus* generalizations or strict ones. If laws cannot explain their instances, then laws that hold only *ceteris paribus* cannot explain their instances. Nor can they explain their non-instances—i.e., what is true when *ceteris* isn't *paribus*. But what is true *ceteris paribus* can tell us about the moral and other dispositions of things, given a suitable construal of '*ceteris paribus*'.

dispositions (tendencies, powers, etc.) of kinds (of actions, factors, etc.), while the latter are the manifestations of such dispositions. And moral generalism, in turn, explains why moral principles can and do play the roles that they play in moral thought, practice, and theory.

Moral dispositionalism also explains the other two theses—besides Holism—with which I began this inquiry. It explains Principle, the thesis that right-making factors or moral reasons entail corresponding moral principles. For if right-making factors are (generic) factors with the power to make actions of a kind *prima facie* obligatory (i.e., dispositionally obligatory overall), and if moral principles ascribe attributes (including powers and other dispositional attributes) to sorts or kinds, then every right-making factor entails a true moral principle—one that ascribes to it the power to make actions of a kind *prima facie* obligatory. Finally, moral dispositionalism explains Conflict, the thesis that conflicts of moral obligation are possible and metaphysically unproblematic. It explains them as concurrent manifestations of conflicting tendencies to be obligatory or wrong—as cases in which either (1) each of two or more mutually incompatible actions manifests a tendency to be obligatory (i.e., to be right in one respect or *prima facie* obligatory) or (2) one action manifests both a tendency to be obligatory and a tendency to be wrong (i.e., to be wrong in one respect or *prima facie* wrong).

This chapter, in part, has been submitted for publication. I am the sole author.

REFERENCES

Anscombe, G.E.M. 1967. Who Is Wronged? *The Oxford Review* 5: 16-17.

Audi, Robert. 1996. Intuitionism, Pluralism, and the Foundations of Ethics. In *Moral Knowledge?: New Readings in Moral Epistemology*, ed. Walter Sinnott-Armstrong and Mark Timmons. New York: Oxford University Press.

_____. 1998. Moderate Intuitionism and the Epistemology of Moral Judgment. *Ethical Theory and Moral Practice* 1: 15-44.

_____. 2002. Prospects for a Value-Based Intuitionism. In *Ethical Intuitionism: Re-Evaluations*, ed. Philip Stratton-Lake. Oxford: Clarendon Press.

Bakhurst, David. 2000. Ethical Particularism in Context. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Bales, R. Eugene. 1971. Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure? *American Philosophical Quarterly* 8: 257-65.

Brandt, Richard B. 1959. *Ethical Theory: The Problems of Normative and Critical Ethics.* Englewood Cliffs, N.J.: Prentice-Hall.

Brink, David O. 1989. *Moral Realism and the Foundations of Ethics.* Cambridge: Cambridge University Press.

_____. 1994a. Common Sense and First Principles in Sidgwick's *Methods. Social Philosophy and Policy* 11: 179-01.

_____. 1994b. Moral Conflict and Its Structure. *The Philosophical Review* 103: 215-47.

Cartwright, Nancy. 1983. *How the Laws of Physics Lie.* Oxford: Oxford University Press, Clarendon Press.

_____. 1989. *Nature's Capacities and Their Measurement.* Oxford: Clarendon Press.

_____. 1999. *The Dappled World: A Study in the Boundaries of Science.* Cambridge: Cambridge University Press.

Crisp, Roger. 2000. Particularizing Particularism. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Dancy, Jonathan. 1981. On Moral Properties. *Mind* 90: 367-85.

_____. 1983. Ethical Particularism and Morally Relevant Properties. *Mind* 92: 530-47.

_____. 1993. *Moral Reasons*. Oxford: Blackwell.

_____. 1999a. Can a Particularist Learn the Difference between Right and Wrong? In *Proceedings of the Twentieth World Congress of Philosophy*, ed. Klaus Brinkmann. Vol. 1: *Ethics*. Bowling Green, Ohio: Philosophy Documentation Center, Bowling Green State University.

_____. 1999b. Defending Particularism. *Metaphilosophy* 30: 25-32.

_____. 2000. The Particularist's Progress. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

_____. 2003. Are There Organic Unities? *Ethics* 113: 629-50.

_____. 2004. *Ethics without Principles*. Oxford: Clarendon Press.

Ewing, A.C. 1959. *Second Thoughts in Moral Philosophy*. New York: MacMillan.

Feldman, Fred. 1995. Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice. *Philosophy and Phenomenological Research* 55: 567-85.

Foot, Philippa. 1983. Moral Realism and Moral Dilemma. *Journal of Philosophy* 80: 379-98. Reprinted in *Moral Dilemmas*, ed. Christopher W. Gowans. New York: Oxford University Press, 1987.

Frankena, William K. 1973. *Ethics*, 2nd ed. Englewood Cliffs, N.J.: Prentice Hall.

Gaut, Berys. 2002. Justifying Moral Pluralism. In *Ethical Intuitionism: Re-Evaluations*, ed. Philip Stratton-Lake. Oxford: Oxford University Press.

Glover, Jonathan. 1977. *Causing Death and Saving Lives*. New York: Penguin.

Hare, R.M. 1963. *Freedom and Reason*. New York: Oxford University Press.

_____. 1981. *Moral Thinking: Its Levels, Method, and Point*. Oxford: Oxford University Press, Clarendon Press.

Hooker, Brad. 2000. Moral Particularism: Wrong and Bad. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

_____. 2002. Intuitions and Moral Theorizing. In *Ethical Intuitionism: Re-Evaluations*, ed. Philip Stratton-Lake. Oxford: Clarendon Press.

Hurka, Thomas. 1996. Monism, Pluralism, and Rational Regret. *Ethics* 106: 555-75.

_____. 1998. Two Kinds of Organic Unity. *The Journal of Ethics* 2: 299-320.

_____. 2001. *Virtue, Vice, and Value*. Oxford: Oxford University Press.

_____. 2003. Moore in the Middle. *Ethics* 113: 599-628.

Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford: Clarendon Press.

Jackson, Frank, Philip Pettit, and Michael Smith. 2000. Ethical Particularism and Patterns. In *Moral Particularism*, ed. Brad Hooker and Margaret Little. Oxford: Clarendon Press.

Kagan, Shelly. 1988. The Additive Fallacy. *Ethics* 99: 5-31.

_____. 1989. *The Limits of Morality*. Oxford: Clarendon Press.

_____. 1998. *Normative Ethics*. Boulder, Colo.: Westview Press.

Kant, Immanuel. 1785. *Groundwork of the Metaphysics of Morals*. In *Practical Philosophy*, ed. Mary J. Gregor. Cambridge: Cambridge University Press, 1996. Prussian Academy pagination.

_____. 1797a. *The Metaphysics of Morals*. Translated by Mary J. Gregor. In *Practical Philosophy*, ed. Mary J. Gregor. Cambridge: Cambridge University Press, 1996 Prussian Academy pagination.

_____. 1797b. *On a Supposed Right to Lie from Philanthropy*. Translated by Mary J. Gregor. In *Practical Philosophy*, ed. Mary J. Gregor. Cambridge: Cambridge University Press, 1996. Prussian Academy pagination.

Kavka, Gregory S. 1979. The Number Should Count. *Philosophical Studies* 36: 285-94.

Korsgaard, Christine M. 1996. Two Distinctions in Goodness. In *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.

Little, Margaret Olivia. 2000. Moral Generalities Revisited. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Lowe, E.J. 1989. *Kinds of Being: A Study of Individuation, Identity and the Logic of Sortal Terms*. Oxford: Basil Blackwell.

McNaughton, David. 1988. *Moral Vision: An Introduction to Ethics*. Oxford: Blackwell.

_____. 1996. An Unconnected Heap of Duties? *Philosophical Quarterly* 46: 433-47. Reprinted in *Ethical Intuitionism: Re-evaluations*, ed. Philip Stratton-Lake. Oxford: Clarendon Press, 2002.

_____. 2000. Intuitionism. In *The Blackwell Guide to Ethical Theory*, ed. Hugh LaFollette. Malden, Mass.: Blackwell.

McNaughton, David and Piers Rawling. 2000. Unprincipled Ethics. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Mill, J.S. 1872. *A System of Logic*. In *Collected Works*. Vol. vii. Toronto: Toronto University Press.

_____. 1873. *Autobiography*. In *Collected Works*. Vol. i. Toronto: Toronto University Press.

Mill, John Stuart. 1871. *Utilitarianism*, 4th ed. Edited by Roger Crisp. Oxford: Oxford University Press, 1998. References are to chapter and paragraph.

Moore, G.E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.

Mumford, Stephen. 1998. *Dispositions*. Oxford: Oxford University Press.

_____. 2000. Normative and Natural Laws. *Philosophy* 75: 265-82.

Nagel, Thomas. 1986. Williams: Resisting Ethical Theory. In *Other Minds: Critical Essays 1969-1994*. New York: Oxford University Press, 1995.

Nussbaum, Martha. 2000. Why Practice Needs Ethical Theory: Particularism, Principle, and Bad Behaviour. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Philips, Michael. 1987. Weighing Moral Reasons. *Mind* 96: 367-75.

Pickard-Cambridge, W.A. 1932. Two Problems About Duty. *Mind* 162: 145-72.

Pietroski, Paul M. 1993. Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory. *Ethics* 103: 489-515.

Prichard, H.A. 1912. Does Moral Philosophy Rest on a Mistake? *Mind* 21: 21-37. Reprinted in *Moral Writings*, ed. Jim MacAdam. Oxford: Oxford University Press, 2002.

_____. 2002. *Moral Writings*. Edited by Jim MacAdam. Oxford: Clarendon Press.

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

Raz, Joseph. 2000. The Truth in Particularism. In *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little. Oxford: Clarendon Press.

Ross, W.D. 1930. *The Right and the Good.* Edited by Philip Stratton-Lake. Oxford: Clarendon Press, 2002.

_____. 1939. *Foundations of Ethics.* Oxford: Clarendon Press.

Sartre, Jean-Paul. 1948. *Existentialism and Humanism.* Translated by Philip Mairet. In *Jean-Paul Sartre: Basic Writings*, ed. Stephen Priest. London: Routledge, 2001.

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

Shafer-Landau, Russ. 1997. Moral Rules. *Ethics* 107: 584-611.

Shope, Robert K. 1965. Prima Facie Duty. *The Journal of Philosophy* 62: 279-87.

Sidgwick, Henry. 1907. *The Methods of Ethics*, 7th. London: Macmillan.

Singer, Peter. 1979. *Practical Ethics.* Cambridge: Cambridge University Press.

Sinnott-Armstrong, Walter. 1999. Some Varieties of Particularism. *Metaphilosophy* 30: 1-12.

Slote, Michael. 2001. *Morals from Motives.* Oxford: Oxford University Press.

Stratton-Lake, Philip. 2002. Introduction to *The Right and the Good*, by W.D. Ross. Oxford: Clarendon Press.

Strawson, P.F. 1949. Ethical Intuitionism. *Philosophy* 29: 23-33.

Taurek, John M. 1977. Should the Numbers Count? *Philosophy and Public Affairs* 6: 293-316.

Tooley, Michael. 1972. Abortion and Infanticide. *Philosophy and Public Affairs* 2: 37-65. Reprinted in *Bioethics: An Anthology*, ed. Helga Kuhse and Peter SInger. Oxford: Blackwell, 1999.

Urmson, J.O. 1975. A Defence of Intuitionism. *Proceedings of the Aristotelian Society* 75: 111-19.

Whiteley, C.H. 1952-3. On Duties. *Proceedings of the Aristotelian Society* 53: 97-104. Reprinted in *Moral Concepts*, ed. Joel Feinberg. Oxford: Oxford University Press, 1969.

Williams, Bernard. 1965. Ethical Consistency. In *Problems of the Self: Philosophical Papers, 1956-1972*. Cambridge: Cambridge University Press, 1973.

_____. 1966. Consistency and Realism. In *Problems of the Self: Philosophical Papers, 1956-1972*. Cambridge: Cambridge University Press, 1973.

_____. 1981. Conflicts of Values. In *Moral Luck: Philosophical Papers 1973-1980*. Cambridge: Cambridge University Press.

_____. 1985. *Ethics and the Limits of Philosophy*. Cambridge: Harvard University Press.

_____. 1988. What Does Intuitionism Imply? In *Making Sense of Humanity and Other Philosophical Papers, 1982-1993*. Cambridge: Cambridge University Press, 1995.

_____. 1995. Acts and Omissions, Doing and Not Doing. In *Virtues and Reasons: Philippa Foot and Moral Theory*, ed. Rosalind Hursthouse, Gavin Lawrence, and Warren Quinn. Oxford: Clarendon Press. Reprinted in *Making Sense of Humanity*. Cambridge: Cambridge University Press, 1995.

Woolhouse, R.S. 1973. Counterfactuals, Dispositions, and Capacities. *Mind* 82: 557-65.