FLORIDA STATE UNIVERSITY

COLLEGE OF ARTS AND SCIENCES

PARTICULARISM AND HOLISM:

NOT A NECESSARY MARRIAGE

By

RICHARD CORDERO

A Dissertation submitted to the
Department of Philosophy
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Degree Awarded:
Spring Semester, 2015

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Richard Cordero defended this dissertation on March 25, 2015.
The members of the supervisory committee were:

David McNaughton

Professor Directing Dissertation

Aline Kalbian

University Representative

J. Piers Rawling

Committee Member

Michael Bishop

Committee Member

The Graduate School has verified and approved the above-named committee members, and certifies that the dissertation has been approved in accordance with university requirements.

For my mother and father.

ACKNOWLEDGMENTS

This dissertation would not have been possible without the incredible support and generosity of a number of people. First and foremost, I would like to thank my family: my parents, to whom this work is dedicated. The brain they did give me has, to a large extent, compensated for the leg they did not. I would also like to thank my sister, Vicky Piscitelli, and my brother Noel Cordero Jr. I could never repay them for the support they gave me over these past ten years, much less the preceeding twenty-two.

My greatest professional debt, of course, is to my dissertation advisor, David McNaughton. Over the past few years, during which this dissertation was written, David proved to be incredibly patient and understanding as I missed self-imposed deadline after self-imposed deadline. David always provided encouragement, even when he had to tell me that I was completely wrong. I consider myself incredibly fortunate to have had the opportunity to study under such a wonderful philosopher, mentor, and man. I would like to thank the members of my committee for their amazing generosity, Mike Bishop, Aline Kalbian, and especially Piers Rawling for his many helpful comments and discussions.

I am also incredibly grateful for the time spent learning from and talking to the many philosophers, friends, and students I have met over the past eleven, or so, years in which I've been engaged in academic philosophy. Among them, Michael Ayoub - for his LaTeX support, Kenneth Barber, William Baumer, James Beebe, Ali Brown - whose high fives are missed, Randy Clarke, Victoria Costa, Russ Dancy, Ian Ellis, Karen Foulke, Josh Gert - one of the most amazing professors in human history, Jack Justus, Stephen Kearns, Hannah Koslowsky, Carolyn Korsmeyer, Shidan Lotfi, Zach Martin, Steve McFarlane, Al Mele, Ben Miller, Jason Miller, Jen Miller, Michael and Lizzie Ruse, John Schwenkler, Josh Shepherd, Ken Shockley - for being the most energetic and passionate professor I have ever had, J. David Smith - for helping me to realize that the answers I was seeking were answers to philosophical questions and not just psychological ones, John Stigall, Pete Takacs, Aron Vadakin, Jamie Watson, Jiyuan Yu, and Chris Zarpentine. I owe a special debt to two friends and colleagues who, sadly, passed away before I had the opportunity to show my appreciation for all they did for me, John Robert "Bobby" Johnson and Heidi Millarker. Bobby helped me to find my second passion, football; my greatest mental and emotional refuge when the work gets too hard

or too frustrating. Heidi helped to renew my faith in the work I was doing at a critical juncture in my studies. They were two amazing people who are deeply missed by all who were lucky enough to have been loved by them.

I would also like to thank two very special friends, Moncho and Rex, who have on a number of occasions, quite literally, saved my life.

And last, but most certainly not ever least, I would like to thank my wife, Gabrielle Graham. Writing a dissertation is a difficult and stressful process. In my case, the difficulty and stress were compounded by my disability and chronic pain. Gabrielle is one of the finest, most caring, most loving, intelligent individuals I could ever hope to have met. Her love and support has helped immeasurably to get me through the final stages of this process. It seems quite impossible to repay her, but I am so grateful that I get to spend the rest of my life trying to do so.

TABLE OF CONTENTS

ABSTRACT

In this dissertation, I examine the prospects for moral particularism. Moral particularism, which, like most views, comes in a variety of flavors, is essentially the view that the role general principles have traditionally played in moral theorizing is overstated. In Chapter One, I lay out the groundwork for the theories which I will discuss in Chapters Two through Four – a framework which I will ultimately reject. The most prominent variety of particularism in the literature, and the subject of Chapter Two, is the one offered by Jonathan Dancy in his *Ethics Without Principles*. (Dancy 2004) Dancy argues for a holistic conception of practical reasons – reasons that sometimes count in favor of Φing can, in other situations, count against Φing. He claims that the truth of holism leads rather immediately (though not directly) to moral particularism. The literature on particularism is flooded with arguments against Dancy. In Chapter Two, I discuss Dancy's particularism and rehearse what I take to be the strongest objections against his view. I take some of those arguments to be rather decisive. And so I suggest that for particularism to survive, we must move beyond Dancy's view.

In Chapter Three I examine another variety of moral particularism - that offered by David McNaughton and Piers Rawling. Given the objections leveled against Dancy, one maneuver for preserving particularism would be to have a variety that is less objectionable by being more like traditional generalist theories. McNaughton and Rawling's particularism is just such a view. However, for various reasons which I discuss in that chapter, I find McNaughton and Rawling's view also to be unsatisfactory.

Given that I take these two, the most prominent, strains of particularism to be the best developed, and given their problems, I suggest in Chapter Four that perhaps if particularism is untenable, particularists might best serve their aims by adopting quasi-generalist views. Essentially what particularists seem to be looking for is a view that privileges moral contexts over moral rules. The rules cannot adequately guide us, not exceptionlessly, at least, given the variety of moral situations in which we find ourselves. So, in this chapter I examine the moral contextualism offered by Margaret Olivia Little and Mark Norris Lance, as well as Pekka Värynen's theory of hedged moral principles. If extant varieties of particularism face insurmountable problems, then perhaps there are varieties

of generalism which privilege moral contexts enough to satisfy particularist intuitions. Ultimately, I find these views also to be inadequate.

The overarching theme of this work is that extant particularist views are inadequate in various ways all relating to their conception of a practical reason. And so, with that in mind, in Chapter Five, I argue that for particularists to maintain a foothold in the debate, they must rethink the conception of a practical reason they employ in their views. I suggest that there is room for a variety of particularism that rejects the traditional conception of practical reasons as holistic contributory considerations that we weigh together and against each other to determine what we ought to do. I call the view I am offering 'Eliminativism', as it is an attempted elimination of contributory-reasons-talk from the discourse. I reject the contributory conception of reasons and offer a more coarse-grained conception of reasons for action that privileges context above all else, thereby giving particularism a coherent conception of reasons for action that enables them to eschew general principles.

CHAPTER 1

THERE ARE REASONS AND THERE ARE REASONS

1.1 Introduction

In this dissertation, I examine moral particularism's prospects. A moral particularist, or particularist, is one who believes that moral agents do not need to rely on moral principles in order to make moral judgments.[1] A moral generalist, or generalist, is one who does believe that moral agents must rely on moral principles in order to make moral judgments. The degree and manner of this reliance is one point of contention between particularists and generalists, resulting in the spectrum of views we find in the literature. For instance, Sean McKeever and Michael Ridge have identified five distinct varieties of particularism. (2006: 15-19) Because particularism is a view about moral principles, the kind of particularist one is can also vary according to the type of principles at issue. The categorization of moral theories as either particularist or generalist is not exhaustive. However, of those who are either generalists or particularists, the vast majority of moral theorists are the former.

One important thing of which we should remain mindful throughout this dissertation is precisely what *kind* of particularism is at issue. As I mention, McKeever and Ridge distinguish between five different varieties of particularism. The distinctions they draw are largely distinctions in the degree of skepticism particularists have toward the possibility or utility of moral principles. However, another way to draw a distinction between varieties of particularism is as follows. One might be a particularist in a metaphysical sense, an epistemological sense, or a psychological sense. That is, one might be skeptical of the very existence of, or work that gets done by, moral principles as parts of the fabric of reality. If we conceive, for the moment, of principles as like laws of nature, then this kind of particularism would be like a skepticism about the existence of the laws or of the causal efficacy of the laws. For instance, such a particularist might claim either that the generalization 'copper is electrically conductive' is either false or not explanatory of particular

[1] I follow Dancy (2004: 7) in particular here in characterizing particularism and generalism.

instances of electrically conductive copper. In the moral domain, a particularist with this sort of metaphysical inclination might claim either that, for example, a moral principle prohibiting the telling of lies is either false or not explanatorily useful in explanations detailing the wrongness of particular lies. On the other hand, one might take on an epistemological flavor to his or her particularism. According to an epistemological strain of particularism, the question is whether or not we require knowledge of moral principles in order to discern moral truths or to make justified moral judgments. Below, in 1.4, I discuss a conception of moral principles according to which grasping a principle is a precondition of conceptual competence with the moral concept cited by the principle. A particularist of this epistemic sort would reject such a conception of principles or any theory that employed them. Finally, the variety of particularism which I think is in fact the most interesting - the psychological conception of particularism. According to this variety of particularism, moral agents as we actually are and actually reason do not rely, or need to rely, on moral principles to make moral judgments and arrive at moral truths. While I think this variety of particularism may be the most interesting, I take it to be largely an empirical position, with largely empirical questions and answers. It is a view about the way in which the brains of moral agents function (the ones we know about). This view is better left to psychologists and cognitive and neuroscientists. In the rest of this dissertation I will *primarily* be discussing metaphysical and epistemological varieties of particularism.

In this chapter, I introduce some of the central concepts and problems that arise in the debate between particularists and generalists. Traditionally one of the central points of contention is just how it is that practical reasons – reasons for actions – function. This is because a common view is that, as David McNaughton has put it, a principle "is...a moral reason which has had its generality made explicit." (1988: 191) It follows, then, that if we figure out how reasons work, we can figure out how principles work. I will sketch a handful of views regarding what reasons and principles are and the various ways in which they work.

1.2 Holism and Atomism

Holism in the theory of reasons is, to put it one way, the view that a consideration, feature, or reason that in one situation supports my performing some action Φ might, in another situation, either not support my Φ-ing or it might count against my Φ-ing. Atomism in the theory of reasons is

a denial of this – it is the view that features carry their normative status with them from context to context. In short, reasons holism is the view that reasons are, perhaps highly, context-dependent. A feature depends on its context for its status as a reason, where 'status' here means not just whether the thing is or is not a reason, but also whether the thing is a reason that counts for or against Φ-ing, and to what degree. Atomism is the view that some reasons are context-independent – the status of some feature as a reason here does not, and perhaps cannot, vary there, in some other context.

For a long while, the debate between particularists and generalists seemed to hinge on the status of the debate between holists and atomists about reasons. Though at least some recognized that there need not be a 1:1 correspondence between particularism and holism, on the one hand, and generalism and atomism, on the other, particularists consistently, persistently, and to varying degrees of strength, made claims to the effect that the truth of holism in the theory of reasons is detrimental to the prospects for a true generalist moral theory. Dancy for instance claims that the holism of reasons is inconsistent with a principle-based approach to moral theory. (2004: 77) David McNaughton has claimed that, according to particularism, "moral principles are at best useless, and at worst a hindrance, in trying to figure out which is the right action. What is required is the correct conception of the particular case in hand, with its unique set of properties." (1988: 190) I take McNaughton here to mean that among the 'set of properties' are the particular reasons that are instantiated in the case at hand. Given this reading, it is attention to the reasons that are present, not to the principles allegedly subsuming the case at hand, which will yield success in moral reasoning.

A common holist tactic is to employ a sort of inductive argument to show that holism is true for one class of reasons (epistemic or aesthetic, for instance), and then to argue that moral reasons are no different except in the fact that they have moral content rather than, say, aesthetic content. That is, we first assume that whatever it is that makes something a normative reason is true of all normative reasons. If that is the case, then normative reasons are all of a piece, and any distinction between, for example, aesthetic reasons and moral reasons must be down to the content of the reason, not to the fact that one is a specifically aesthetic reason while the other is a specifically moral one. For example, suppose R_1 is a an aesthetic reason that favors appreciating comedian A over comedian B. Suppose that comedian A has better timing. Suppose R_2 is a moral reason

that favors helping someone who is in need. The important differences between R_1 and R_2 for the purposes of the holist's arguments, are due to the fact that R_1 is about comedic timing and R_2 is about the distress of another person. The argument would then proceed by providing an example in which R_1 is a reason that, either does not favor or counts against, appreciating one comedian over another one. Suppose that the comedian is telling horribly racist jokes immaculately. The fact that her timing is impeccable, that she waits the perfect amount of time between the setup, the punchline, and the button, or that her callbacks are all flawlessly placed in her set, seem as though they all make the comedy worse.[2] Holist arguments sensibly start with simple and obvious cases displaying the apparent context-dependence of reasons and then move to the more contentious cases. What this does is it puts the burden of showing that the different classes of reasons discussed are in fact different kinds of reasons squarely on the shoulders of the atomists.

Another example illustrating holism comes from Dancy. Suppose I have before me a pill that will make blue objects appear red and red objects appear blue. Normally, when I see a blue object, the fact that I am seeing a blue object is a reason for me to believe that there is a blue object before me. However, after I have taken the pill, the fact that I see a blue object is no longer a reason for me to believe that there is a blue object before me. Instead, this fact is a reason for me to believe that there is a red object before me. While of course this does not show that all epistemic reasons behave holistically, it does show that not all of them behave atomistically. (2004: 74) What about other types of reasons? Suppose Alice has just been invited to a dance party. Alice is a big fan of dancing. That there will be dancing at this party is a reason for Alice to go. A day before the party, however, Alice's favorite uncle, Frank, dies of a heart attack, sending Alice into period of deep grief. Given this grief, she will no longer be able to enjoy dancing at the party. That there will be dancing is no longer a reason for Alice to go to the party. Again, this does not show that all practical reasons can change with context, but it is evidence for the view.

Particularists rely on the claim, which strikes me as reasonable, that if practical reasons generally behave holistically, then moral reasons do as well. There would have to be something importantly different about specifically moral reasons to justify the view that, despite being, in the main, a variety of practical reasons, they function differently than other practical reasons. This view strikes me as rather unlikely. Perhaps an example of a moral reason changing with context would

[2]These claims, of course, rest on the assumption that one should not appreciate horribly racist humor. I take this point as uncontroversial.

help show that it is, at least, possible, and that therefore it is not the case that all moral reasons behave atomistically. So here is a well-worn example of alleged moral reason-switching. Usually, the fact that some act would be the telling of a lie is a reason not to do it. If I am with my legally blind uncle at the clothing store helping him shop for a suit for a wedding to which he is attending, I should answer him truthfully when he asks if the colors of the items he has selected match. It would be wrong to lie to him about whether a lilac shirt matches neon green pants. However, there are other situations in which it seems that it would be wrong not to lie – situations in which the fact that something is a lie does not speak, even a little bit, against it. Suppose a dying friend, deeply in the throes of insanity due to her illness, asks me to see to the safety of her stable of unicorns after she passes away. I can think of no reason not to lie here in this case. I tell my friend that it would be my pleasure to take care of her unicorns for her. The fact that this is a lie seems to me not to count at all against it. In fact, to the extent telling this lie eases my friend's mind in her last few moments of life, it seems the right thing to do. Take a different case. Suppose I were kidnapped and forced to share a meal with a despot, after which he is giving a speech that will be televised all over the world. During the course of the meal, a piece of food gets stuck in his teeth. As we are finishing up, he stands and asks, "How do I look?" I know that if he gives his speech with a large piece of kale stuck between his teeth, he will look very silly and it will embolden his critics; it will serve to undermine his, presumably, illegitimate authority. It seems to me that I positively have a reason to lie to him. Given the despotic nature of his government, anything that makes his critics braver can only be a good thing. Here, the fact that my telling him, "You look great!" is a lie, does not seem to count at all against my so telling him.

The preceding examples are, of course, not meant as proof of the holism of all reasons. What they are meant to do is to show that at least some reasons behave holistically. The question now is, if these reasons behave holistically, why not others? It seems that any alleged atomistic feature, consideration, or reason, must either be a different kind of reason, and so be a member of a different kind of thing from the features and reasons employed in the preceding examples, or it must be atomistic, not in virtue of the fact that it is a reason, but rather because of the particular content of the reason it is. (Dancy 2004: 77) For example, the most difficult examples for particularists to offer showing reasons changing their valence, or, the direction of their support, seem to be those involving the thick moral concepts – justice, courage, temperance, etc. – and

if these features behave atomistically, it must be because of something about these thick moral concepts themselves that make the reasons involving them behave atomistically. As Dancy notes, if there are these invariant reasons, this is not a problem for holism as a theory of how reasons work, since reasons qua reasons are context-dependent (if holism is true). This is consistent with there being *some* reasons whose contents are special in that they do not admit of variability. (2004: 77) Either way, the burden is on the atomist (or perhaps quasi-atomist if she admits that some reasons behave atomistically while others behave holistically) to show how and why these reasons and not others behave atomistically if she insists that they do.

1.3 Reasons

Reasons are tricky things. One important question to which I will not give much attention is whether or not there really are reasons. That is, is there room in a respectable ontology for reasons? It seems in our everyday lives, we talk about and offer what we take to be reasons pretty regularly. We talk about and offer reasons probably at least as frequently as we talk about what we want for dinner or whether our favorite sports team won their latest match (in my case, they usually have not). The point here is just that whether they are real or not, whether we are mistaken about them or not, talking about reasons plays a very important role in our lives. I will take it for granted that reasons do exist. I do not believe anything hinges on exactly what their mode of existence is, for instance whether they are real or merely a useful fiction. In the rest of this section, I will examine a few of what I take to be the best candidate accounts in the literature of what reasons are if they do exist in some form or another.

However, before doing that I should like to make a few remarks about just what kind of reasons are the relevant type here. Unless otherwise indicated, throughout this dissertation I will be discussing normative practical reasons. These are to be distinguished from other kinds of reasons – the kind of causal/explanatory reason that we are looking for when, say, we ask why the espresso machine is not working. Is it plugged in? Does it need water? Did we push the wrong button? The kind of reason that would be offered here in answer to these three questions would be a purely

descriptive fact that explains why the espresso machine did not work. The kind of reason of interest herein is not this causal/explanatory kind.[3]

Normative reasons somehow or other are the kinds of things that give support to, or count in favor of, the performance of some action. One thing to note is that it is not always clear in the literature whether what particularists and generalists are saying about moral reasons is meant to hold true for all types of reasons or not. Now, one's views about reasons, even one's view about moral reasons, need not commit one to either particularism or generalism.[4] However, some views do seem to lend themselves better to either one view or the other. At the very least, certain proponents of certain views claim that their view of what reasons are and how they work lends itself better to either particularism or generalism. For instance, Dancy's *Ethics Without Principles* offers a careful examination of how thinking of reasons in the way he does leads to moral particularism. (2004)

Once we narrow our focus, for the sake of the debate between moral particularists and moral generalists to just practical (moral) reasons, we encounter still other issues. Exactly what various particularists and generalists take reasons to be and do, and whether the conceptions they hold are even coherent, is not always clear. While there are a number of different avenues to take here, the most sensible is surely to look at just what it is that arch-particularist Jonathan Dancy thinks about practical reasons.

1.3.1 Dancy on Reasons

Dancy (2006; 2004) has written in several places that while it may appear as though there are a number of different things reasons might do, the majority of those views fail in at least one important respect, whether it is because they simply fail to capture what we think reasons are or because they beg the question against one or more of the components of Dancy's view. It is thus via proposing a criterion that his own view sufficiently meets and showing how other views fail to do so that Dancy establishes a beachhead from which to build his particularist position.

When Dancy talks about reasons, he is talking about what he, and many others, call 'contributory reasons.' A contributory reason is a feature of a situation that makes something of a case

[3]There is one minor caveat, however. Below I will be discussing John Broome's view of reasons as explanations, however, the view of reasons-as-explanations in Broome's sense is distinct from reasons-as-explanations in the espresso machine case.

[4]For instance, people who hold very similar views of reasons might disagree about whether particularism or generalism, on the one hand, or atomism or holism, on the other, are true. This is the case, in fact, with many particularists and generalists. They all agree that reasons are contributory, but disagree about whether that entails anything about the truth of particularism and holism.

for performing some action. (2004: 15) They count in favor of, or against, the performance of actions. Because they stand in this relation between action and agent (i.e. they count in favor of some *agent* performing some *action*), Dancy calls reasons 'favorers' when they count in favor and 'disfavorers' when they count against. The reason they are called contributory reasons is that they merely contribute to the case for or against some action without settling the matter. That is, the way in which contributory reasons count in favor of performing actions does not foreclose the possibility that the case for performing an action can be made better or worse by the introduction or removal of other features in the situation. These other features might themselves be contributory reasons or they might be any of several types of other features that Dancy countenances which need not themselves be contributory reasons. Dancy calls these other features enablers, disablers, intensifiers, and attenuators. (2004)

A feature E is an enabler if its presence in a situation enables another feature R to be a reason for action. Similarly, a feature D is a disabler if its presence in a situation disables another feature R from being a reason for an action. Enablers turn features that might not have otherwise been reasons into reasons while disablers prevent features that might have otherwise been reasons from being reasons. Intensifiers and attenuators are similarly two sides of a coin. The presence of an intensifier N in a situation raises the strength, as it were, of some reason R. Similarly, the presence of an attenuator A in a situation decreases the strength of some reason R. Intensifiers and attenuators are not the kinds of features that can change the status of another feature from reason to non-reason, but they are not themselves wholly irrelevant in a situation. Perhaps the role these features can play will be clearer with a pair of examples. (Dancy 2004: 38ff.)

If I borrow a book from my friend Anne and make an uncoerced promise to return that book, my having promised is here a favorer – it carries normative force and thereby supports my returning the book. The fact that the promise was uncoerced is an enabler. In the absence of this fact, if my promise had been coerced, then my promising would not favor my returning the book. Here, the uncoerced nature of the promise counts as a local enabler – one specific to this case. There are also global enablers. The fact that I am able to return the book is one such enabler, because if something like a 'reasons implies can' principle is true (whereby it cannot be true that I have a reason to Φ unless I am able to Φ), then there is always at least one enabler present whenever I have a reason to do anything. Suppose further that I know that my friend borrowed the book

from the university library, that I promised to return it by noon today, and that it is due back by the end of the day today. The fact that the book is due back by the end of the day today is neither a favorer itself, nor is it an enabler – turning the promise from a non-favorer into a favorer. Instead, this consideration is what Dancy calls an intensifier – it strengthens the case for returning the book when I promised to return it. On the other hand, suppose as I am preparing to leave my house to return the book I begin to get a migraine headache. I could still return the book, but the intensity of the migraine makes it difficult for me to gather the strength to lift myself off of my couch. Anne knows that I am prone to migraine headaches, is understanding about it, and the fee from the library would be no more than whatever change happens to be lying between the cushions of my couch. Even though I could still return the book, it seems as though the fact that I have a migraine acts as an attenuator on the strength of my promise, diminishing the strength of it as a favorer – not disabling it or overriding it – just making the demand to fulfill the promise a little less stringent.

Dancy thinks that the various forms of relevance considerations can take, that is, either being a favorer or disfavorer, an enabler or a disabler, or an intensifier or attenuator, lends itself very well to particularism without begging any serious questions against generalism. This is because, in the first place, thinking of reasons this way is easily, but not necessarily, construed as lending itself to the holism of reasons. Considerations can be, but need not be, "turned on" (made favorers by enablers) or "turned off" depending on what other considerations are involved in a given situation. If this is so, then some consideration R can, in one situation, favor A-ing, while in another situation it can either not favor A-ing or perhaps even favor not-A-ing. This would be because in other situations there might be a disabler present. Dancy believes that, while holism does not entail particularism, its truth makes particularism much more likely than generalism. One reason to think this is the fact that according to holism, details matter – details, even minor ones, can change the roles played by various normative features across contexts. Particularism is the view that this variability resists codification – that no finite list of finite principles can capture all there is to capture about morality. While it was at one time largely believed (or more accurately, simply assumed) that the truth of holism led almost straightaway to particularism, various examples have been given that purport to show that generalism is compatible with holism as well. [5]

[5]See, for instance, McKeever and Ridge's (2006) example involving a simple Utilitarian theory (U). I discuss (U) in the next chapter.

Dancy thinks drawing the distinctions he does between the three forms of relevance is the best way to make sense of the way contributory reasons function. One serious, and to my mind perhaps legitimate, criticism against this view Dancy offers is that it seems to lack parsimony. What Dancy is offering is a metaphysical view explaining how the normative landscape is carved up. There are bits here that count as reasons, distinct bits here that are like reasons, but are not reasons, but are the things on which reasons depend for their status as reasons. There are still other bits of the landscape that can stand in important relations to reasons and enablers, but are not either, and what these bits do is they intensify or attenuate the strength of a reason. Dancy has seemingly gone feature-individuation crazy. Why should we think that intensifiers and enablers are distinct, rather than being part of the overall ground of a reason? It might seem much simpler to posit a long background condition upon which the reason carrying the day in a particular case depends, rather than to posit myriad distinct entities each pulling the reason this way and that.

On the one hand, a good rule of thumb seems to be that we should not posit more entities in our ontology than there need to be. On the other, though, sometimes things are messy. Sometimes things seem complicated when they are simple and sometimes they seem simple when they are complicated. Anyone who has seen an episode of Scooby Doo should be familiar with such cases. The existence of ghosts, monsters, and pirates could, in a sense, easily explain almost all of the villainy that went on in the various towns Scooby and his gang visited. However, admitting ghosts and monsters into our ontology complicates matters in a different way than the way in which the true nature of the various crimes committed complicates matters. The various forms of relevance that Dancy says features of a situation can take without question is a messier way of looking at normativity than other views. However, unless these other views can capture all that we need them to capture, and better than Dancy's view, we may not have any reason to prefer one position over the other.

1.4 Principles

Just as there are various ways to think about what reasons are, there are also various way to think about what principles are. Often, in the literature, particularists and generalists alike switch seamlessly and effortlessly between talking in terms of principles and talking in terms of reasons. If McNaughton's remark is right about principles being reasons who have had their generality made

explicit, then this would account for the connection between the two, and thereby for the frequency with which theorists slide from one to the other. In fact, drawing the connection tightly this way between reasons and principles, i.e. viewing reasons as having a built-in generality can, as I will discuss in Chapter 2.2.2 and 2.2.3, lead to problems for particularists. The idea here is meant to be that if we know that some reason Φ favors performing some action here, then the next time we come across Φ, it will favor performing the same sort of action and this just is, roughly, what it is for something to be a principle of action. Whereas the relationship between a reason and a principle is, for most generalists, a very tight one, for most particularists these two notions (must) come apart. The debate, then, is whether and how reasons are keyed into principles. If they are, then what sort of principles do we get? While particularists are most happy to talk about reasons all day long, we tend to become somewhat squeamish whenever any substantive principles are brought into the discussion.

One way of carving up the terrain is the way Dancy (2013) does it. According to Dancy, we can have two kinds of principle: absolute or contributory. Absolute principles are universal claims that all actions of a particular type are wrong. For instance, 'all lies are wrong' is an absolute principle. Such principles are problematic for a variety of reasons. One such reason is that, given the complexity of morality and our everyday lives, we are very likely to come across situations in which more than one principle applies. When that happens, all of the applicable principles must be in agreement – they must all say that our action is either right or that it is wrong. If they disagree, then we are faced with an action that is both absolutely right and absolutely wrong. We should want to avoid such contradictions in our moral theory. Therefore, we should reject the idea of an absolute moral principle.

The other conception of a moral principle which Dancy (2013) discusses is that of a contributory principle. Like his conception of a contributory *reason*, a contributory principle identifies some feature or set of features as contributing to the rightness or wrongness of an action. Actions which have those features are, to that extent either right or wrong. Whether an action is overall right or wrong is determined by the overall balance of rightness and wrongness articulated in the relevant principles. Throughout this dissertation, when I speak of principles, I will generally be speaking of contributory principles. Instances in which that is not the case will, I hope, be clear given the context.

Before discussing any concrete examples of alleged moral principles, we had better get a bit clearer on just what a moral principle is supposed to be as well as on what a principle is supposed to do. There are many things that a principle can do. McKeever and Ridge (2006) identify six criteria for what makes a moral generalization count as a moral principle.[6] I take each of these criteria to be sufficient, but not necessary for a generalization's being a principle. The six criteria are as follows. The generalization must be a standard, in the sense that it provides the application conditions for moral concepts. One caveat here, though, is that the application conditions cannot be trivial. That is, the truth of supervenience, of the moral on the non-moral, already entails the existence of exceptionless application conditions for moral concepts. But so-called 'supervenience functions' are not at all the kind of thing we are looking for in a moral principle. They are unwieldy – too long to be of any use to normal moral agents. Secondly, the application conditions provided ought to provide an explanation of why the moral concept in play applies when it does. So principles qua standards must provide *non-trivial* application conditions explaining why the moral concept applies when it does. A fully specified principle articulating the wrongness of lying, for instance, would provide different explanations depending on the theory of which the principle was a part. For Kantians, for example, the principle might cite some inability to universalize the maxim countenancing lying. (McKeever and Ridge 2006: 7-8)

A second criterion is that the generalization must provide guidance. That is, it must be useful in practice. As McKeever and Ridge (2006: 8) note, a given standard might be theoretically useful, but practically useless. In particular, if the application conditions for some moral concept are extremely complex, we might never be able to employ that principle. So, we should like our principles also to provide some practical guidance. We can combine the 'standards' condition with the 'guidance' condition to yield a third criterion, that generalizations must act as action-guiding standards in order to have the status of moral principles. (McKeever and Ridge 2006: 9)

Interestingly, the fourth criterion that McKeever and Ridge discuss yields a sort of principle that virtually no theorist actually employs. It is that a generalization must provide an algorithmic decision-procedure. That is, like a mathematical formula, if we have all of the inputs and know what the relevant operators mean, we should have to do very little thinking about what the result is.

[6]The distinction between a generalization and a principle is a rather subtle one. The thought is that not just any generalization qualifies as a moral principle. "Any action is either morally required or it is not" is true, and it is a generalization, but it falls far short of qualifying as a moral principle. (McKeever and Ridge 2006: 5-6)

(2006: 11) However, this criterion has a serious problem. In order to apply the principle, one must first have the judgment necessary to identify one's situation as being subsumed under a particular principle. But this is just the sort of judgment from which the algorithmic nature of the principle is meant to free us. It is no wonder that virtually no one holds this criterion as the standard for what makes a mere generalization a moral principle.

The fifth criterion claims that moral principles are preconditions of competence with moral concepts. (McKeever and Ridge 2006: 11-12) The thought is that in order to understand whether, for instance, a particular lie is permissible or impermissible, one must be in possession of the relevant principles. Not having a grasp of the principle entails not understanding lying sufficiently well enough to make moral judgments regarding its rightness or wrongness in particular cases. This variety of moral principle seems to be one of the central targets of Dancy (2004), wherein he characterizes particularism as the view that the possibility of moral thought and judgment does not depend upon a suitable supply of moral principles.

Finally, the sixth criterion McKeever and Ridge discuss is that generalizations act as truth-makers for particular moral truths. McKeever and Ridge usefully cash out this criterion with an example. The generalization 'all the coins in my pocket are made of copper' is the sort of generalization which is true, if it is, only accidentally. Had I paid the store clerk differently than I did, I might have a different set of coins in my pocket. So this generalization, if true, is *made true* by the particular fact that all of the coins in my pocket happen to be made of copper. On the other hand, consider the law of nature 'copper is electrically conductive'. This generalization is different than the one about the coins in my pocket. Rather than being made true by all of the instances of copper's electrical conductivity, it is the instances of copper's electrical conductivity that rely on the truth of the generalization. The thought, then, is that moral principles are like 'copper is electrically conductive' rather than like 'all the coins in my pocket are made of copper'; they are like laws of nature – the preconditions for particular truths. (2006: 12-14) That said, however, McKeever and Ridge note that this conception of a moral principle is rather orthogonal to the debate between particularists and generalists. I present it here just to provide a completed account of what McKeever and Ridge claim are the ways in which generalizations get to have the status of moral principles.

In addition to the above criteria, there have traditionally been three desiderata of moral principles. The candidate principle must be true, useful, and exceptionless.[7] First, obviously, the principle must be true. There might be some room to give here, depending on how one conceives of morality, though. For instance, non-realists about morality might insist that principles are necessary and useful, but not true. Such fictionalists might defend moral principles on pragmatic grounds.[8] But setting aside these and other skeptical views, that a principle is true seems to me the most important criterion.

Second, the principle must be useful. A principle that is, for example, either too general or too specific would not be useful. Candidates that are too general, e.g. 'do the right thing', 'do what's best', 'do what the virtuous person would do', or 'do what you have most reason to do', are not useful, though they are obviously true. On the other hand, the candidate principle also must not be too specific. A principle which cites a particular agent's name or a particular time and place would be useless in all but one instance. Instead, a more plausible candidate might say something like, 'Help those in need when you are able to, if doing so will not cause you or anyone else significant harm.'

A third desiderata, which has in recent years been the subject of much discussion,[9] is that the candidate principle must be exceptionless. One tactic for holding on to principles used to be to add modifications whenever a successful counter-example was posed. If we had the principle, 'Do not lie' but were then faced with a counter-example involving the choice between telling a lie and saving someone's life, a defender of the candidate principle might modify it to include just this sort of exception. We would then have something like, 'Do not lie unless it is to save a life.' We could, of course, come up with several other plausible counter-examples involving the choice between telling a lie and something other than the loss of someone's life but more serious than telling a lie, such that the defender of the principle might continue to expand the principle to account for all of the exceptions we pose. Ultimately, however, continuous expansion to head off counter-examples will leave us with a principle so long that it would no longer be useful. We might have something like, 'Do not lie unless it is to save someone's life, or unless it is a trivial matter and the choice is between lying and sparing your significant other's feelings, or unless it is about something trivial and the

[7] As I will discuss shortly, this last criterion is contentious.
[8] See, for instance, Joyce (2001), esp. Chs. 7 and 8.
[9] See, for instance, Lance and Little (2007) and Väyrynen (2006; 2008; 2009).

truth will cost you your house, or unless...' and so on. One way to avoid having such lengthy principles would be to include normative content in the principle. In Chapter Three, I discuss David McNaughton and Piers Rawling's preferred brand of particularism, which admits of certain evaluative principles. The problem for generalists with these sort of already-evaluative-principles is that part of what principles are meant to do is to bridge the gap between the non-normative and the normative. That is, they are supposed to help us to identify which non-normative features in the various situations we come across are normatively relevant.

In an effort to get around the problem of the ever-expanding principle, some theorists (generalists and particularists alike) have taken a cue from principles we find in the sciences, particularly biology, employing ceteris paribus, or hedged, principles.[10] Hedged principles purport somehow to have the exceptions built-in, and therefore do not have to specify exceptions. That 'fish eggs turn into fish', a stock example from the literature, is an example of a hedged principle. The interesting thing about this principle is that, as it turns out, in the vast majority of cases, fish eggs do not turn into fish. Rather, in the vast majority of cases, fish eggs become food for other creatures before they have the opportunity to hatch. Nonetheless, we still think of it as true that fish eggs turn into fish. As Lance and Little note, what the principle 'Fish eggs turn into fish' does - the reason why it is an acceptable principle is that it tells us something important about the nature of fish eggs. (2008: 61) It seems there would be something faulty about our knowledge of fish and their eggs if we did not believe that fish eggs turn into fish. If proponents of hedged principles are onto something significant, then it seems that the criterion that any candidate principle be exceptionless is not a genuine criterion – it is too stringent. As long as the candidate principle works the way we want a principle to work, then it need not be exceptionless.

1.5 Moving Forward

In the rest of this dissertation, I examine the two most prominent varieties of moral particularism. In Chapter 2, I discuss Jonathan Dancy's particularism, which McKeever and Ridge (2006: 19) dub 'Anti-Transcendental Particularism.' I argue that there are challenges to Dancy's view which warrant our looking elsewhere for a viable particularist candidate. In Chapter 3, I turn my attention to the view offered by David McNaughton and Piers Rawling. Their view is explicitly

[10]See for instance Väyrynen (2006; 2008; 2009) and Lance and Little (2006a; 2006b; 2007; 2008) and Little (2000)

more amenable to a host of generalist intuitions. Given the failure of Dancy's view, however, and the plausibility of certain particularist tenets, I think it is worth looking at whether some extant variety of particularism can stand in. Ultimately, the generalist concessions that McNaughton and Rawling make are both problematic and too anti-particularist for someone looking for a strongly particularist view. However, if both of the best extant particularist views are problematic, then perhaps particularism itself is irresolvably flawed. For that reason, in Chapter 4 I examine whether two generalist views, the views offered by Mark Norris Lance and Margaret Olivia Little and the view advanced by Pekka Väyrynen. The idea is that, if particularism really does have all of the problems that I will discuss, then perhaps particularists need to seek a generalist view which can accommodate their most important intuitions. Because Little (2000) is an expression of a variety of moral particularism and because Lance and Little (2006a; 2006b; 2007; 2008) are all the articulations of a view which agrees with particularists about the status of a certain variety of moral principle, there should be something in the view that can still appeal to particularists, despite the generalist elements. Similarly, Pekka Väyrynen's view takes as one of its starting points the truth of the holism of practical reasons. If Dancy (2004) is right that thinking about reasons holism leads almost straightforwardly to moral particularism, then Väyrynen's view must surely not be *too* objectionable for particularists. I argue, however, that the moves Väyrynen makes are unacceptable and, at points, beg important questions against particularists. Therefore, finally, in Chapter 5 I offer an account of reasons that I think helps particularists to salvage their view, while sacrificing only minor aspects of the traditional particularist framework.

1.5.1 Particularist Enough?

Throughout this dissertation I will be criticizing a variety of views, both generalist and particularist. Some of the criticisms I raise rely, as I indicate above, on the position that the views in question are "not quite particularist enough," or not strongly particularist. I take this position in Chapters 2 through 4, while criticizing certain moves Dancy makes (Chapter 2), the view offered by McNaughton and Rawling (Chapter 3), and in my assessment of Lance and Little's moral contextualism as well as Pekka Väyrynen's theory of hedged principles (Chapter 4).[11] Here I would

[11]Strictly speaking, simply in virtue of the fact that neither Lance and Little nor Väyrynen are offering particularist views, their views fail to meet the 'particularist enough' requirement. As I hope to make clear in the course of the discussion of the views, there are adequate reasons for examining their views in the context in which I examine them.

like to make a few brief remarks as to what I mean by 'particularist enough' and why this is not an unreasonable or unfair requirement of the views I discuss.

There are, I think, one or maybe two basic requirements a view must meet in order to be 'particularist enough'. First, it would seem odd for a particularist view to require an understanding of, or competence with, a moral principle in order for an agent to be able to apprehend a moral situation appropriately. It would be odd for a particularist view to require that an agent have a grasp of a moral principle as a precondition for the agent's actually being a morally competent agent. This would make a generalist conception of morality a precondition for moral agency. This would mean that no particularist view could possibly be correct. So any view with this is a feature would not be 'particularist enough'.

And so, secondly, or perhaps as a corollary of the first requirement, any appearance of moral principles in a view must be downstream, as it were. That is, the theory should not start with objectionable moral principles. If principles appear in the theory, they should appear as bridge principles, not from non-normative features to normative ones, but rather from normative features to other normative features. That is, these principles should tell us, not what reasons we have given the non-normative features of a situation, but rather overall what we ought to do given the reasons that we have. Any principles from the non-normative level of a situation to the normative level (the level at which reasons appear), seems necessarily to be context-insensitive in a way that is contrary to a central particularist tenet. Particularism, as I conceive of it, is in the business of drawing our attention to the context-sensitivity of the moral features of the world. Given the myriad ways in which the states of affairs we come across every day can vary, any principle trying to bridge this first level normative gap would have either to be so broad as to capture irrelevant features of the situation (or to leave out relevant features of the situation) or it would be too complicated for use by human moral agents. But once reasons are in play, there might not be anything objectionable about principles bridging the gap between the reasons we have and what we ought to do.

As a first pass, these seem perfectly legitimate desiderata of a particularist view. They are merely articulations of specific questions that might be begged against particularism. While we all beg some questions sometimes, at some level, when engaged in a discourse, we should try to beg as few questions as possible. In the rest of this dissertation, I will try to show how and why the views I discuss in Chapters Two through Four either beg important questions against particularism or are

problematic in other ways. Ultimately, in Chapter Five, I hope to provide a variety of particularism that satisfies the most important particularist intuitions.

CHAPTER 2

THE BASIC PARTICULARIST ARGUMENT

2.1 Introduction

Arguably, no single philosopher has done more to bring attention to the potential problems for moral generalism over the last thirty years than Jonathan Dancy. Having sketched Dancy's conception of contributory reasons as context-dependent favorers in Chapter One, in this chapter I will examine in more detail Dancy's broader particularist view. In the first section, I describe the basic particularist argument - a general statement of the particularist position and argument that can be read from Dancy's work. In the second section, I canvas some objections to Dancy's view as well as, when possible, offer replies on Dancy's behalf. Finally, I offer some remarks to the effect that, while Dancy's particularism may be correct in the main, his particular version of it has flaws that justify looking elsewhere for a different variety of particularism.

2.2 The Basic Particularist Argument

Jonathan Dancy, as most particularists have, claims that the truth of reasons holism bodes ill for the prospects of a generalist morality. (Dancy 2004: 77) That is, if reasons are as context-dependent as Dancy believes them to be, then we are left without any considerations to which we can reliably attach general moral principles. If the fact that some action of mine will cause either myself or someone else a significant amount of pain sometimes counts against my performing that action and sometimes counts in favor of my performing that action, then the fact that some action of mine will cause either myself or someone else a significant amount of pain cannot reliably indicate whether such actions ought to be prescribed or proscribed. Particularists offer myriad examples involving alleged morally relevant considerations operating in non-standard ways. For instance, sometimes the fact that someone is in need of help is a reason to offer that help. If I come across someone bleeding in the street, their distress seems clearly to be a reason for me to stop and call an ambulance. Sometimes, however, the fact that someone is in need of help is not a reason to offer that help - as in the case of a child who gives up on a homework problem much too quickly.

In the first case, offering help would clearly be a good thing, whereas in the latter case, one might not unreasonably think that helping the child too soon would somehow be doing the child harm or a disservice.

There are a few different ways one can construct the argument against generalism. Here I reproduce a line of argument that we can find in Dancy (2004), which I think effectively captures the general particularist position.

> P1. The holism of theoretical reasons is uncontentious.
>
> P2. Theoretical and practical reasons are of the same basic kind; the logic of reasons is the same in both domains.
>
> P3. The holism of practical reasons is true.
>
> P4. The logic of moral reasons is the same as that of practical reasons generally.
>
> P5. Moral reasons function holistically.
>
> P6. Moral principles specify features as general, or atomistic, reasons.
>
> P7. If moral reasons function holistically, then the possibility of such reasons cannot rest on the existence of principles that specify morally relevant features as functioning atomistically.[1] (Or, if the possibility of moral thought and judgment depends on a suitable supply of principles, then moral reasons must function atomistically.)

> C. The possibility of moral reasons does not depend on the existence of principles that specify morally relevant features as functioning atomistically. (Or, the possibility of moral thought and judgment does not depend on a suitable supply of principles.)

Support for P1 comes from a catalog of available cases like the following. Sometimes the fact that the ground is wet is a reason to believe that it has recently rained and sometimes it is no such reason. When the morning meteorologist claims a 70% chance of rain, this gives credence to the belief that the ground is wet because of rain. However, on a hot summer's day during a draught, the ground's being wet is more likely due to something else - perhaps the opening of a fire hydrant. I mentioned Dancy's own example in the previous chapter, in which I have taken a pill that makes blue objects appear red and red objects appear blue. While under these effects, arguably, the fact that I am perceiving a blue object is not a reason for me to believe that there is a blue object

[1] Dancy (2000: 135)

in front of me. Rather, it is a reason for me to believe that there is a red object in front of me. Normally, I am not under such effects, so seeing a blue object is a reason for me to believe there is a blue object in front of me. All this is basically to say that the apparent evidence can sometimes be misleading. What we take to be a reason for some belief here might be a reason going the other direction in another situation.

Support for P2 is perhaps somewhat more complicated and more difficult to provide. I noted in the previous chapter that the claim that theoretical and practical reasons are of the same basic kind, that there is one logic of reasons, is a default assumption. Without an argument to the effect that there is a fundamental difference between theoretical and practical reasons, other than the contents of the particular reasons themselves, it is reasonable to believe them to be basically the same type of thing.

As luck would have it, though, there are a few arguments that, if good, seem to show that there are important differences between theoretical and practical reasons. If there are such differences, then it may not be the case that there is just one logic for all reasons. If there is no such single logic, then it seems reasonable to suspect the move from claims about theoretical reasons to claims about practical reasons. I will consider this and other objections in the final section of this chapter. For continuity's sake, I will discuss the rest of the argument for the general particularist position first.

The third premise is the claim that the holism of practical reasons is true. This premise follows fairly straightforwardly from P1 and P2 - P2, of course, being something for which further argument must be offered.

The fourth premise is the claim that there is just one logic of, at least, practical reasons. Whether or not theoretical and practical reasons function in the same way, it seems very plausible that practical reasons all function in the same way. Moral reasons are just a particular type of practical reason, and in whatever way practical reasons work, moral reasons should function in basically the same fashion. If the average, mundane, non-moral practical reason functions atomistically, then probably, so does the average moral practical reason.

The fifth premise follows fairly straightforwardly from P3 and P4, while P6 is a fairly simple claim about how a certain type of moral principle works. Here is where much recent debate between particularists and generalist has taken place. In recent years, generalists have moved away

from the claim that morality can be codified by mostly exceptionless principles. Now, and much more prevalently, generalists claim that the type of principles that form morality's framework are principles that admit of exceptions: hedged or ceteris paribus principles. Whether this is, or could be, the case, and how the possibility of useful, exception-ridden principles bears on the debate between particularists and generalists will be the subject of the latter half of Chapter 4, in which I discuss Pekka Väyrynen's theory of hedged principles. Väyrynen's hedged principles are all the more important to the generalist position if Dancy's argument against exceptionless principles is sound.

The seventh and final premise is the claim that, assuming what generalists need is exceptionless principles, these principles cannot make use of atomistic reasons if what morality consists of is reasons that function holistically. If moral reasons function holistically, then any principles employing those reasons could not represent those reasons as functioning atomistically. This would be something like trying to measure the weight of a stone without a constant gravitational pull. The stone has a particular weight on the Earth, a different weight on the moon, and a different weight for all the various possible gravities to which it might be subject. If we don't know beforehand which gravitational constant to use in our calculations, we cannot predict the weight of the stone. Similarly, if we do not know the valence of the reason in play (if the holism of reasons prevents us from knowing beforehand how the reason will combine with the other features of the situation), then we cannot know beforehand whether a given principle will apply in the situation. This, as we will see throughout this dissertation, is at the heart of the particularist position - the supremely complex nature of moral situations and the context-sensitivity of reasons.

The conclusion, then, of the Basic Particularist Argument is just that if reasons function holistically, no principle can employ atomistic reasons and at the same time be exceptionless and true. One may have, of course, a variety of reasons for rejecting the BPA as it has been presented here. In the rest of this chapter I will discuss a handful of objections raised against particularism.

2.3 Objections to the BPA

The following objections are what I take to be some of the strongest objections to particularism that are in the literature. Because some of these objections are so different from the others, the things a proponent of the BPA might say in response to one objection might preclude a certain type

of response to another objection. Ultimately, it seems to me this might be what happens and is a reason for jettisoning at least certain aspects of the BPA. For, if multiple objections are severely damaging to the particularist position, but one of the responses I offer precludes the possibility of a satisfactory response to a different objection, then the choice for the proponent of the BPA is about merely which way the argument fails.

Before approaching the more difficult objections, we might do well to look at one obvious potential response to particularism. One might wonder whether, if what particularism needs to get off the ground is holism, why not just reject the holism of practical reasons? So, one way might be to insist that reasons, or at least (some) moral reasons, do not function holistically. If this were the case, then the generalist could still employ some exceptionless principles tracking moral reasons that are context invariant. For instance, many theorists wish to draw a distinction between what we might call 'basic', or underived reasons, and derivative reasons. Such theorists pushing the anti-holistic line of argument might claim that basic reasons are atomistic while derivative reasons are context-dependent, and to that extent holistic. One thing the anti-holists or anti-particularists might say in this context is that the fundamental moral principle(s), whose existence is evidence of particularism's falsity, depend on the existence of basic reasons, rather than derivative or holistic reasons. This line of argument, however, seems somewhat intuitively implausible to me. In the first place, it seems that there are at least some considerations that are strong candidates for basic reasons which do not seem to function atomistically. For example, many people think that pain is bad. Pain, for such people, functions as a basic reason. However, it seems plausible that pain might, in different contexts, both be a reason for and a reason against acting in certain ways. That Φ-ing will cause S pain can be a reason for S to Φ if S believes in the utility of self-flagellation. Cases in which pain acts as a reason against acting are obvious. Of course, it is open to the anti-particularist to claim that whatever basic reasons we cite as being context-dependent are merely apparently basic reasons - they are in fact derivative reasons. The fact that we find a reason that is context-dependent does not mean that there are no invariant basic reasons, it just means that we have not picked out a basic reason at all.

At this point in the dialectic, it seems we have come to an impasse. The particularist thumps the table claiming that an alleged basic reason can change its valence in a particular context - it is holistic. The anti-particularist then thumps the table and claims that that alleged basic reason is,

in fact, a derivative reason. So the particularist has done nothing to harm the anti-particularist's claim that there are still some basic, atomistic reasons from which to derive our moral principles. There are, it seems two ways to approach the problem now.[2] Either we stay mired in the trench warfare that is the example counter-example tactic, or we press ahead and try to show that there is something about reasons, in principle, which tells for or against the particularist or anti-particularist position. The prospects of this latter tactic will largely depend on that theory of reasons that we are working with. As we will see throughout this dissertation, that is a truly important question. For now, however, and for the rest of this chapter, I turn to other, hopefully less impasse-inducing objections.

2.3.1 McKeever's and Ridge's Objection

Perhaps a more promising avenue to take against the basic particularist argument is to claim that, though reasons function holistically, this does not mean we cannot generate principles that are consistent with the holism of reasons. In fact, several generalists and most particularists admit that the holism of reasons is *not* inconsistent with generalism. For instance, Sean McKeever and Michael Ridge (2006) have argued that there are ways to articulate a form of utilitarian principle that employs holistic reasons. They offer the following principle, which they call (U):

> (U) The fact that an action would promote pleasure is a reason to perform the action if and only if the pleasure is non-sadistic. The fact that an action would promote pain is a reason not to perform the action. An action is morally right just in case it promotes at least as great a balance of reason-giving pleasures over pain as any of the available alternatives; otherwise it is wrong. (2006: 29)

Since (U) is a complete codification of the morality of right and wrong in purely descriptive terms, it is inconsistent with particularism (2006: 29). However, McKeever and Ridge take (U) also to presuppose holism about reasons. (U) picks out pleasure as a context-dependent feature - providing reasons to perform actions when the pleasure is non-sadistic and not otherwise. Since (U) is inconsistent with particularism and yet presupposes holism about reasons, it cannot be the case that holism provides direct, or even unique, support for particularism. If McKeever and Ridge are right here and if particularists require unique support from holism to get their view off the

[2]This is a point Dancy (2004: 2) makes at the outset of his book.

ground, then they are in serious trouble unless they can respond. Fortunately, there is perhaps a response to be made on behalf of the particularist.

First, it is not true that particularism requires unique support from the holism of reasons to get off the ground. If the argument is that there is no reason to suppose particularism is true if generalism is consistent with the holism of reasons, then it is a bad one. Granted, certain particularists claimed that the fact that holism is true bodes ill for generalist morality.[3] That this is false, though, has little to do with whether or not generalism is true. As I understand it, the truth of holism has no *direct* link to the truth of either particularism or generalism.

One reason for assessing the merits of particularism, independent of whether or not generalism and holism are consistent, is the fact that generalists themselves are stuck in debates about how correctly to codify morality. Utilitarians, deontologists, virtue ethicists, and every other kind of generalist all have at least a handful of hard cases for which his or her view seems to provide an inadequate or counter-intuitive answer. Perhaps the reason why generalists struggle to find a consistent set of finite principles that codify morality is that there is no such set. By way of a more direct response to the concerns raised by (U), however, particularists might wonder about the way McKeever and Ridge are individuating the relevant considerations.

One way to read (U) is as McKeever and Ridge have laid it out - pleasure is sometimes a reason and sometimes, when it is sadistic, it is not a reason. Therefore, pleasure's status as a reason is context-dependent. However, we might wonder whether it is in fact *pleasure* that provides reasons or whether, perhaps instead, it might be the fact that the pleasure is non-sadistic. McKeever and Ridge insist "it is the fact that an action would promote pleasure that is a reason when it is a reason and not the fact that it would promote non-sadistic pleasure." (2006: 29) One might think, as I do, that there is still something McKeever and Ridge ought to say about just how or why it is pleasure and not non-sadistic pleasure that counts as a reason. Their insistence amounts merely to table-thumping. It seems somewhat an open question, *pace* McKeever and Ridge, whether, if pleasure is alleged sometimes to be a reason and sometimes not to be one, it isn't the non-sadistic nature of the pleasure or the pleasure itself that is doing the work. Here it seems that McKeever and Ridge have left themselves open to a line of argument that similar to one they raise against Dancy

[3]As, for instance, when Dancy (2004: 77) says, "A principle-based approach to ethics is inconsistent with the holism of reasons."

(2004) in the context of his discussion about default reasons.[4] Without getting too far afield, they complain that Dancy has failed to distinguish between the absence of a condition which prevents a default reason from being the reason it is and the presence of a condition which enables that same consideration to be the reason it is. So, for example, when making a promise, if the promise was made under coercion, except for in very special circumstances, we tend to think that the fact of coercion is a defeater for the bindingness of the promise; it is a disabling condition. The question is, then, what is the reason? Is it the fact that a promise is freely made or is it the fact that there was a promise, together with a separate condition - the promise's having been freely made? These questions about reason-individuation are a tricky point that seem not to get too much attention in the particularism and generalism literature. Likewise, is the right way to read (U) to take pleasure or non-sadistic pleasure as the relevant consideration? That is, because there are no cases (given (U)) in which pleasure that is not non-sadistic is a reason for someone to act, on what grounds can the generalist identify the pleasure itself as a reason rather than the compound feature of non-sadistic pleasure? According to (U), the property an action has of producing pleasure counts as a reason to perform that action only when that pleasure takes a non-sadistic mode. Without an explanation, there seems no good reason to put (U) the way they do. Likewise, without that explanation, what reason do we have for supposing that pleasure is functioning holistically in (U) rather than for supposing that non-sadistic pleasure is functioning atomistically? I think, then, that McKeever and Ridge have not done enough here to show that generalism is consistent with holism. While, I suspect that holism and generalism are consistent, as I will discuss in Chapter 4 when I discuss Pekka Väyrynen's theory of hedged principles, I do not think it is a threat to particularism.

Whether McKeever and Ridge have a damning argument against the particularist or not may not matter, however. Joshua Gert (2007) and Selim Berker (2007) both, though independently, have argued, in effect, that whether generalists can articulate principles employing a holistic conception of reasons is irrelevant. Particularists, in particular those following Dancy have a larger problem. The way particularists employ holism, according to Gert and Berker leaves them without a coherent notion of a reason.

[4]See McKeever and Ridge (2006: 46ff)

2.3.2 Berker's Argument Against Particularism

In "Particular Reasons," Selim Berker describes particularists as adhering to a set of views, the conjunction of which leads the particularist to a serious problem, viz. that of finding herself with an incoherent conception of a reason for action. (2007) According to particularists, at least according to those particularists who adopt Dancy's framework, the way reasons function is holistic through and through. That is, there is a holism at the non-normative level in the way non-normative features come together to form contributory reasons. Holism appears here at the contributory level as well, in the way that contributory reasons come together to yield what we might call overall reasons, moral verdicts, or the moral status of an action. So we have three levels - non-normative, contributory reasons, and overall moral. At the non-normative level, we have the sorts of things that are described using non-normative terminology - pleasure, pain, punch, etc. At the contributory level we have the way in which the elements at the non-normative level combine to form contributory reasons. For example, we might have some person's pain giving me a reason to X we have the fact that this is a false claim giving me a reason to Y, we have the fact that Z-ing will cause someone some non-negligible amount of harm giving me a reason to do it or not to do it, etc. Finally, at the highest level, we have the many ways in which all of our X-ings, our Y-ings, and our Z-ings come together to yield some overall verdict - some account of what I have most reason to do, or as we might put it, the overall moral status of the action we are about to perform. Sometimes, it might be right to flip the switch, causing the harm to some while giving the pleasure to others, and at other times it might be wrong to do so. According to the holism of reasons, the way the constituents of reasons combine to yield reasons depends on all (relevant) elements of the situation. However, this leads precisely to the problem Berker is raising.

That is, problems arise for particularists when we examine more closely just how elements at the contributory level combine to yield overall reasons. According to the normative framework within which particularists all seem to be building their view, the sorts of things reasons are, are the sorts of things that can combine to create a better or worse case for whatever it is those reasons speak for or against. For example, the fact that I promised a friend to meet her for lunch tomorrow is a reason for me to go to lunch. The fact that I promised and the fact that I haven't seen this friend in a very long time seems an even stronger reason to meet her for lunch tomorrow. The fact that I promised and the fact that I haven't seen her in a very long time combined with the fact that I

had to cancel our last three attempts at having lunch seems still a stronger reason. The fact that reasons can combine like this to create stronger reasons in some cases and in other cases weaker ones, seems to require an explanation; it seems as though we ought to be able to explain according to what function reasons combine - we need a combinatorial function (Berker 2007: 120). Whatever this combinatorial function is, it cannot be, for the particularist, a simple arithmetical function; certainly not simply an additive function. If it were simply an additive function, then it would be very difficult to see how particularism could be true. After all, if all we have to do is add the weights of our reasons, then we can read principles straight off of the various sums we might get. Therefore, part of the particularist theory of reasons includes what Berker calls Noncombinatorialism about Reasons for Action (NCR). NCR is the view that the combinatorial function for rightness and wrongness is not finitely expressible, and so in particular, is not an additive one. (Berker 2007: 119-22) However, according to Berker, this view leaves us without a coherent notion of a reason at all.

Berker attempts to illustrate his objection by way of an analogy involving a simplified version of Newtonian Classical Mechanics. The following scenario is meant to be analogous to the view against which particularists argue. It is meant to be a 'generalist' version of Newtonian classical mechanics, if you will. He asks us to suppose we have a series of particles in some section of space interacting with each other only through the classical laws of gravitation. Each particle exerts a force on every other particle that is a function of relative positions of the particles and their masses. We can measure the individual force of one particle on another and that force will stay constant regardless of the positions or masses of any of the *other* particles in the space around our two particles. Indeed, if our two particles were to move, so long as their positions relative to each other remain the same, the individual force one particle exerts on the other will stay the same as well. We can also calculate the total force acting on a particle by taking into account the masses and relative positions of all of the particles in this section of space and calculating the force each other particle exerts on the particle in question. The fact that the individual force exerted by one particle on another will remain constant insofar as their relative positions remain the same, no matter when or where the two particles are, is analogous to atomism about reasons, according to Berker. In the way the individual force is insensitive to the positions of any other particles or anything else going on in surrounding space, so too reasons maintain their valence

across contexts, irrespective of any other considerations appearing in any new contexts. The total force exerted on a particle is meant to be analogous to the claim that particularists deny because of their commitment to NCR. We can determine the total force exerted on a particular particle if we have all of the relevant numbers (masses and positions) just as we can, by aggregating all weights of the contributory reasons features presented to us in a situation, determine what the overall moral status of an action is. (Berker 2007: 122-23)

A particularistic version of this Newtonian model would look just as one might expect it to. What we might wish to call the individual force exerted by one particle on another is a result both of the properties held by and relations between those two particles as well as the properties held by and relations between those two particles and all of the other particles in the system.[5] Since according to a particularistic version of the Newtonian model there is both a holist aspect and a noncombinatorial aspect of the system, the relative positions of two particles can be held fixed, while the other elements in the system change, and this can result in a different individual force exerted on one of the fixed particles by the other held-fixed particle. So, the individual force acting on one particle due to another is not given by any general formula that holds regardless of the positions and masses of the other particles. (Berker 2007: 123)

Of course, when we see things this way, we see that what we are calling the individual force is not an individual force at all, really. The force exerted on one particle by another is a *function* of all of the forces and particles in the context - there is nothing individual about it. What *could* we mean by individual force here? Berker proposes a few conceptions of 'individual force', none of which seem to work. It cannot mean that if no other forces were acting on the second particle, it would accelerate in the direction exerted by the force of the first particle, since the nature of the force the first particle exerts on the second is determined, at least in part, by the other forces and particles in the context. For the same reason it cannot mean that if the individual force were absent, then the particle would accelerate in the opposite direction. Nor can 'individual force' pick out the contribution to the total force acted on one particle by another particle. This is because given the noncombinatorialist element of the particularist framework, the total force acted on one particle is not always a simple additive function of the forces acting on it by other particles. It is therefore difficult to make sense of the contribution that one particle can make to the total force

[5]Here I say, "What we might wish to call the individual force..." because, if Berker is right, given this particularistic framework, we do not have a coherent conception of an individual force.

when the way in which that total force is achieved is unspecifiable. So, Berker wonders, to what can the notion of an individual force being exerted on one particle by another really amount?

If the Newtonian model is analogous to the model of reasons employed by particularists, then we can ask similar questions about what the notion of a reason for action is supposed to mean. To show the conception of a reason for action is lost to the particularist, Berker discusses conceptions of a reason for action that are analogous to the conceptions of individual force already discussed and rejected. These conceptions of a reason for action are all fairly clearly rebuffed by Dancy in several places throughout the literature.[6] One common conception, which we might call the isolationist conception, is that a reason for action is a consideration that would, if it were the only relevant consideration, count decisively in favor of performing that action and a reason against action is a consideration that would, again, if it were the only relevant consideration, count decisively against performing that that action.[7] There are at least two problems with this conception of a reason for action. The first problem is that the account seems perhaps to be circular. Talking about the 'relevant consideration' is, in this context, tantamount to talking about the very reason in question. If we are trying to give an analysis of a reason for action then it is no use talking about relevant considerations - the relevant considerations are the ones of which we are trying to give an analysis, viz. they are the reasons.[8] A second problem with this conception of a reason for (or against) action is that if what makes some consideration a reason for action is the fact that it would carry the day if it were the only relevant consideration in some similar but counter-factual situation, then every consideration that might ever potentially be a reason is a reason always. On this account, any consideration carries the day if it is the only consideration. This, however, is false. Not every consideration is a reason in every situation - even if we might wish to say that for every possible consideration there is some possible situation in which it is a reason.[9]

A second conception of a reason for action that Berker considers and rejects on behalf of particularists is, in a way, the reverse of the isolationist conception. We might call this the removal conception of a reason for action. According to the removal conception of a reason for action, a reason for action is a consideration which, were it removed from a situation, its removal would result in a weaker overall reason (for or against). We can reject this conception on particularist

[6]See, for instance Dancy (2006; 2004)
[7]See, for instance, (Ross 1930: 19-20)
[8]Dancy (2006: 41) raises this sort of criticism of Ross's conception of a prima facie duty.
[9]See Dancy (2006) and Dancy (2004: ch.2-4) for further criticisms of this and other conceptions of reasons for action.

grounds in the following way: According to this conception of a reason for action it cannot be the case that the removal of a reason for acting can strengthen the overall reason for performing the action for which the removed reason was a reason. But of course, holists may wish to say just this sort of case is possible. Indeed, Dancy has a rather plausible example showing why we should reject this removal conception. Suppose I am considering performing some action that would benefit a friend. The fact that it would benefit my friend is good and is a reason for me to perform the action. However, suppose that in an exactly analogous case the person for whom I can perform the action is not a friend. It seems possible that in some situations like this the action would be better or I would have more reason to perform the action if the person for whom I'm performing it is not my friend. (Dancy 2006: 43; Berker 2007: 126) Arguably, almost any case of open-handedness or generosity would be a case like this. I could give this money to a friend who asked for it, and that would be good, or I could give this money to a hungry, homeless person who asked for it, and that might be better. The point is just that the particularist, in virtue of her commitment to holism, wants to leave such a possibility open and the removal conception of a reason for action closes this possibility.

A third conception of a reason for action, which Berker calls the 'right-making conception' is the view that a reason for action is a consideration that counts in favor of some action being the right thing to do. So, if here the fact that my action will cause some person harm is a reason against performing that action, then the fact that my action will cause some person harm counts negatively toward my action being the right thing to do (or, it counts positively toward my action being the wrong thing to do). Berker claims that the particularist cannot take this conception of a reason for action on board because, absent an additive combinatorial function, this conception does not make much sense. Berker goes on to provide a formalized explanation of just what he means by this. The point, I think, can be summed up as follows: we cannot make sense of a consideration combining with another consideration to result in a better or worse case for some course of action if we cannot make clear sense of how they can combine. Unless the combinatorial function is additive, or at least quasi-additive, the function by which the reasons combine might result in a reason which, though individually counting in favor of action, it actually decreases the strength of the overall case for action (or, we might have a reason against action strengthening the case in favor of action). (Berker 2007: 127-28)

The final conception of a practical reason that Berker considers is that a reason for action is a consideration that counts in favor of action and a reason against action is a consideration that counts against action.[10] The distinction here between this conception, which we might, following Berker, call the 'favoring conception' and the 'right-making conception' is that the latter takes practical reasons to count in favor of their actions' being the right thing to do, whereas the favoring conception counts in favor of actions full-stop. Whether these two conceptions are super distinct or not does not really matter - as the reason, according to Berker, that particularists cannot take on board the right-making conception applies to the favoring conception as well. Talk of reasons combining in ways that increase the degree to which an action is favored by the set of reasons available in a given situation only makes sense when the function according to which these reasons are combined is an additive function. [11]But if the first two conceptions of a reason for action must be rejected on holist grounds and the latter two must be rejected because they conflict with the particularist's commitment to noncombinatorialism about reasons, what is the conception of a reason for action with which particularists are working?

Dancy, of course, takes contributory reasons to be favorers - he is committed to the favoring conception of reasons for action. However, if Berker is right, Dancy must either reformulate his conception of a reason for action or he must reject his commitment to noncombinatorialism. Berker notes, it seems to me rightly, that the conflict for particularists can be resolved but that in many ways the cost of resolving this conflict - between their commitments to reasons holism, noncombinatorialism about reasons, and the generalized weighing framework - would be too high for many particularists. (Berker (2007: 133ff.) There seems to me to be something very intuitive about holism about reasons. Likewise, there is something very appealing about the generalized weighing

[10]This conception of a practical reason is the one adopted by, for instance, Raz (1990) and Scanlon (1998).

[11]Berker does grant that the combinatorial function might not be strictly additive but quasi-additive. Rather than the function being a straight calculation of the weight of one reason plus the weight of another reason, the function might, so long as the valences of the reasons remain constant, weight certain reasons differently. For instance, if we have only two relevant considerations, one of which should be weighted (for whatever reason) more than the other, we can represent the combinatorial function as follows:

$$R_1 + (R_2)^3 = R_n$$

So we have the weight of the first reason combined with the weight of the second reason cubed giving us the total weight of the reason in favor or against performing some action. Here it matters whether the second reason is cubed or merely squared since, if the second reason is a reason against (and so represented by a negative number), cubing it will maintain the polarity of the reason while squaring it will result in the number representing its weight being a positive number (i.e. reason in favor).

framework. It might seem as though the least unintuitive move for particularists to make would be to give up noncombinatorialism about reasons - but this is precisely what makes them particularists. Generalists, it seems, can and do take on board both holism and the generalized weighing framework - Ross (1930) might be one example of such a generalist, Scanlon (1998) is another. While there might be other avenues of response, [12]I suspect that particularists may ultimately have to reject the generalized weighing model of reasons, and perhaps also the holism of reasons, for action in order to make their case against generalism. In making these moves, particularists will have to take on board an entirely different conception of a reason for action. I will explore a possibility along these lines in Chapter 5.

Now I turn to two objections to particularism, or aspects of the particularist position, both of which were developed by Joshua Gert. The first of Gert's objections is very much like the Berker objection in that the argument put forth aims to show that the particularist hasn't got a coherent concept of a reason for action.

2.3.3 Gert's 1st Argument Against Particularism

According to the generalized weighing framework, normative reasons have one role to play in practical deliberation. This role, however, is rarely carefully and explicitly articulated. Phrases like, 'practical reasons favor action' and 'practical rationality requires that one act on the strongest reason', appear frequently in the literature, but further details are only infrequently offered. In a series of papers, Joshua Gert argues, in effect, that while there is something very intuitive about speaking of reasons in the terms used by proponents of the generalized weighing framework, there is at the same time something very misguided about it. The way proponents of the generalized weighing model speak of reasons makes it seem very much as though reasons play a single role and

[12]See, for instance, Lechler (2012), in which she argues that there is a mistake in Berker's reasoning. Berker assumes that particularists cannot, because of holism, know the quasi-additive combinatorial function by which reasons combine. Lechler argues that, if particularist know sub-functions, which might be quasi-additive, they might thereby know the quasi-additive combinatorial function by which reasons combine. It is unclear to me, though, whether such a combinatorial function would reflect reality in any way. By that I mean, it is unclear whether agents do or could reason in this way. It is for this reason that I do not discuss Lechler's arguments at greater length. The argument relies on what appears to me to be an overly intellectualized account of how agents might reason. A normative view, in particular a view of reasoning, is only worth anything if it reflects how the relevant agents either do or can reason.

This, in fact, is why, despite my earlier protestations against discussing what I called 'psychological particularism', I will go on in the final chapter to discuss the way in which agents might actually reason (as opposed to how they *ought* to reason.)

have only one dimension of strength. Given the intuitive pull of the weighing metaphor to which reasons theorists often appeal, it should not be surprising that the manner in which they claim reasons are weighed against one another is as simple and straightforward as it is.[13]

However, as Gert notes, there seems to be more to the role reasons play - a clear logical distinction between two roles in particular - that cannot be accounted for by the generalized weighing framework. One role reasons play is a justifying role. There are certain actions that are usually irrational. These actions, though, might be made rationally permissible given certain considerations - certain reasons. Those considerations act as justifying reasons for the performance of that action - they help to explain its rational permissibility. To borrow Gert's example, normally running into traffic would be irrational. However, if one runs into traffic to save a child's life, then one's action is made at least rationally permissible by the fact that one is running into traffic to save a child's life. (Gert 2007: 537) The other role Gert claims reasons play is the requiring role. Some actions might be irrational given certain considerations, though normally they are rationally permissible. When these considerations are present they act as requiring reasons - requiring the non-performance of the action in question. For example, in normal situations, it is rationally permissible for me to take a pain-reliever for joint pain. However, if I have recently been prescribed some other medication that would react violently with a pain-reliever, then I am rationally required not to take the pain-reliever. Here, the fact that the medication would react violently with the pain-reliever acts as a requiring reason - requiring me not to take the pain-reliever. (Gert 2007: 538-39)

At first blush one might think that because actions can be made irrational or rationally permissible given the appearance of certain considerations, that Gert's view might lend itself to the particularist project. One might think that the reason that an action's rational status changes is because of changes in the reasons an agent has with respect to that action. However, Gert believes that reasons' strength values must remain stable in order to have a coherent conception of reasons at all. Like Berker (2007), Gert (2007) raises an argument against the particularist to the effect that her commitment to holism leaves particularism with something of a nonsensical conception of the strength of a reason for action.

According to Gert when we talk about the strength of a reason, we are unavoidably implying certain conditionals about that reason. In effect, a reason must have a stable strength value,

[13]See in particular Gert (2007).

otherwise it is unclear what it is that is being influenced by the other considerations in the case. This is because, "[a] strength value is a way of representing regularities of a certain kind: it is an index that tells us where, in an ordered series of conditional claims, truth leaves off and falsity begins." (Gert 2007: 554) Similarly, according to Gert, when we talk about the strength of a reason, we are committing ourselves to certain conditionals about the way it affects the rational status of an action in comparison with other reasons across a range of cases. So, for example, Gert claims that when one assigns a greater degree of justifying strength to the fact that some action will save a stranger's life than one assigns to the fact that some action will save one's own finger, one is engaging in one way of endorsing a claim like the following: if a sacrifice is rationally permissible for the sake of saving one's own finger then that sacrifice is likewise rationally permissible for the sake of saving a stranger's life, but not vice versa. This claim, of course, is a conditional claim about the sorts of relative justifying strengths certain actions may have with respect to certain sacrifices. (Gert 2007: 553-56) The problem particularists face here can perhaps be better illustrated via an analogical example Gert discusses involving the strength or power of a fishing line. Suppose we have three fishing lines each with a different strength-value. Suppose we have a 10-pound line, a 15-pound line, and a 20-pound line. Now suppose we attach a 17-pound weight to the end of each line. What it means for these fishing lines to have the strength ratings they have is that the 10-pound line and the 15-pound line will break, while the 20-pound line will hold. The 20-pound line holds because it has a particular strength rating; a stable strength rating indicated by our being able to call it a 20-pound line. If, when we attached a 17-pound weight to a 20-pound fishing line, sometimes the line held and sometimes the line broke, it would be difficult to understand what we could possibly mean by calling that line a 20-pound line. (Gert 2007: 553)

The problem alleged for particularists is that if they want to claim that the rational (and perhaps moral)[14] status of an action is a result of the strengths of the reasons favoring and disfavoring that action, then they need some account of what the strength value of a reason indicates. Gert claims that particularists have no such account if the strength value of a reason in a given context fails to indicate a stable pattern of strength values across a range of contexts. For Gert, what explains

[14]If the generalized weighing framework is indeed the framework within which particularists are working - and it clearly does seem to be, at least in Dancy's case - then the arguments Gert makes against particularism in practical reason seem as though they apply equally well to moral particularism. This is because, despite the possibility of distinguishing the realms of practical rationality and morality, they have the same structure. In both domains, at least given the generalized weighing framework, the overall rational/moral status of an action is the product of the combination of whichever reasons favor and disfavor it.

the strength of a reason in a particular case is tied inextricably to the way that reason functions in other cases. One tempting maneuver a particularist might make at this point is to claim that they do not need stability. It is enough to get a sense of the strength of a reason to see that reason working in context. It is, after all, part of the particularist project to privilege context and to encourage a move away from attempting to assess situations before all of the details of the context are manifest. Gert, however, attempts to preempt this avenue of response by claiming that if he is right, particularists have no account of the strength of a reason *even if* they try to claim that a particular reason has a particular strength in a particular context. Indeed, Gert claims, "it is completely unclear what [the claim that a reason has a particular strength in a particular context] could mean unless it means that it can be assigned a value that could be used to figure out the rational status of other actions in other contexts." (Gert 2007: 556) It is my hope that by the end of this dissertation, in particular in Chapter 5, I will have provided a response to this maneuver by Gert. Ultimately, the response entails leaving behind the generalized weighing framework, thereby avoiding the issues that Gert and Berker have raised.

2.3.4 Gert's 2nd Argument Against Particularism

Moral particularists often point to an alleged analogy between theoretical reasons and practical reasons in an effort to motivate their view. The idea is that, if theoretical reasons and practical reasons function in roughly the same way or according to the same logic, then the truth of theoretical reasons holism is evidence for the truth of practical reasons holism. This is the maneuver made in the basic particularist argument, premises 1-3. Dancy claims it would be incredible if the logic of moral reasons were so drastically different from other reasons such that moral reasons functioned atomistically while other reasons functioned holistically. He then discharges his now familiar red/blue objects example as evidence of theoretical reasons holism and suggests that practical reasons holism is likewise true. (Dancy 2000: 132) Mark Lance and Margaret Olivia Little likewise appeal to the analogy in their "Defeasibility and the Normative Grasp of Context". (Lance and Little 2004: 436ff.)[15] The thought seems to be, roughly, that reasons have the same basic logical structure and behave in the same basic way across domains. That is, epistemic reasons behave the same basic way that moral reasons do, the same basic way that aesthetic reasons do, the same

[15]It is worth noting that Lance and Little are no longer happy to be called particularists - now preferring the label 'moral contextualists' - because they eschew the (majority of) particularists' insistence on rejecting all generalizations.

basic way that prudential reasons do, etc. So, if epistemic reasons behave holistically, then so do practical reasons.

As I discussed in Section 2.2.1, McKeever and Ridge have, in effect, suggested that the truth of this analogy offers no assistance to particularism because practical reasons holism is not evidence of the truth of particularism. Gert argues against particularism on roughly the same point. His argument, however, is not that practical reasons holism fails to offer any assistance to particularism. Instead, Gert argues against the employment of the analogy itself. In "Putting Particularism in its Place," Gert claims that the analogy between theoretical reasons and practical reasons does not hold. (Gert 2008: 312-324)

In both the epistemic domain and the practical domain, there are two related notions. In the epistemic domain, we have the truth of a proposition and the rationality of believing a proposition. In the practical domain, we have the objective rationality of an action and the subjective rationality of an action. In the epistemic domain, reasons are related more directly to epistemic rationality than to truth. Despite what a straight analogy between the epistemic domain and the practical domain may lead us to believe, Gert says, practical reasons are more directly related to objective rationality than to subjective rationality. (2008: 312)

Gert thinks that the reason why we should not be surprised that particularism[16] is true in the epistemic domain is that epistemic reasons are more directly related to what it is epistemically rational to believe and what it is epistemically rational for a human to believe is determined, at least in part, by the "oddly-contoured limitations" of human nature. Because the human cognitive apparatus is basically the result of a sequence of historical accidents, we should expect epistemic rationality to include some messiness - some variability. (2008: 315-16) The story in the practical domain, however, is somewhat different. If Gert is right up to this point, then, he claims, we should not expect the analogy between epistemic reasons holism and practical reasons holism to hold.

The objective rational status of an action is a matter of whether an action of that type 'makes sense.' If we understand objective rationality in this way, then, Gert claims, there are some plausible candidates for generalist-friendly reasons that favor or oppose actions. Some examples of such reasons are that an action would increase the risk that someone will suffer pain, death, or injury.

[16]Gert uses the term 'particularism' throughout, though he rarely talks at all about the codifiability of reasons or principles. It seems to me that Gert may do better to use 'holism', as the topic of discussion at hand seems more clearly to be the way reasons function – atomistically or holistically. I will, from this point forward, use 'holism' rather than 'particularism.'

These reasons count as reasons independently of whether anyone has epistemic access to them (just as a truth is true independent of whether anyone has epistemic access to it). (2008: 315-16)

Given these allegedly invariant considerations, here is where the analogy begins to break down. The problem for holists is that whereas in the epistemic and practical domains truth and objective rationality are analogous notions, reasons do not bear on truth while reasons *do* bear on the objective rational status of an action. Reasons do not bear on truth for two reasons: 1) reasons can be for and against the things for which they are reasons, and 2) it does not make sense to speak of reasons against a particular *truth*. (2008: 315) Now, practical reasons *can* and do bear on the objective rational status of an action. We can make sense of increasing or decreasing the strength of practical reasons - for instance by increasing the risk or degree of harm suffered. But, Gert claims, it does not make sense to talk about increasing or decreasing the strength of epistemic reasons, at least, to the extent the relevant notion is truth and not the rational status of holding a belief. But the latter notion is analogous to the subjective rational status of an action in the practical domain. So, the way reasons come to bear on the relevant notions in the epistemic and practical domains is different.

One thing some particularists may be willing to do here is to resist that they have to adopt this weighing model for fine-grained considerations-as-reasons in either domain. Selim Berker suggests that this is a potential way out of the worry that he raises for particularists. (Berker 2007: 134ff.) This move would have particularists abandon dependence upon individual considerations and instead focus on the larger context as a whole. However, Berker also thinks that this would generate a new set of problems that are just as difficult as the one the maneuver was meant to solve. Likewise, such a maneuver made in an effort to handle Gert's arguments would likely prove to generate a new set of problems, also just as difficult to deal with. Again, in Chapter 5 of this work, I attempt to get around these concerns by taking steps away from the generalized weighing framework which seems at the heart of these problems for particularists.

Gert's position depends on the distinction between objective and subjective rationality that he draws and their alleged analogs in the epistemic domain. Gert tells us that the distinction between objective and subjective rationality depends on the following three ideas: (1) Objective practical rationality exists, (2) the objective rational status of an action is a function of the relevant practical reasons, and (3) someone can fail to act objectively rational in some instance, and yet

be undeserving of any criticism in terms of his or her subjective practical rationality, because of, say, justifiable ignorance of the relevant reasons. Gert contends that there are analogues of (1) and (3) in the epistemic domain, but not of (2). So the analogy apparently breaks down. The analogs are: (1') there is such a thing as truth, (3'), someone can fail to believe some truth and yet be undeserving of any criticism in terms of his or her epistemic rationality because he or she may have been justifiably ignorant of the relevant facts. The reason why there is no analog of (2) is because, truth-value is not a function of the relevant epistemic reasons. Truth, again, is not determined by reasons, since reasons can count for and against, but there is no sense to be made of talk about reasons for or against a truth. All we get are reasons in favor or against *belief* in some truth. (2008: 319-320)

A potential problem for Gert may be that the alleged distinction between subjective practical rationality and objective practical rationality really is not a substantive one. Recall, Gert says that the objective rational status of an action is a matter of whether an action 'makes sense.' Subjective practical rationality is "a matter of an action's being produced by mechanisms that are as reliably productive of objectively rational action as can be expected of beings with our idiosyncratic limitations..." (2008: 313) But what exactly does this amount to? It seems true that a human action cannot, in any significant and true sense of the phrase, 'make sense' independent of human interests. If this is right, then it is difficult to see how the supposedly objective rational status of an action could avoid being infected by the messiness that goes along with human nature, i.e. the messiness that is a product of our idiosyncratic limitations, i.e. the messiness that renders actions subjectively rational or irrational. So both objective and subjective rational status are, it seems, inextricably tied to to the contingencies of our human lives.

Even if this were true, however, Gert claims that his argument would go through anyway because it would still be the case that the relation between objective and subjective rationality would be different than the relation between truth and epistemic rationality. (2008: 319) However, it is not clear to me that Gert tells us what he takes the relationship between truth and epistemic rationality to be. He tells us that the role of objective practical rationality in the practical domain is the same as the role of truth in the epistemic domain (though he does not tell us precisely what he takes this role to be). (2008: 313) He tells us that reasons are more directly related to epistemic rationality, that reasons cannot determine truth, that we cannot have reasons for or against a truth, but he

does not seem to tell us just what the relationship is between epistemic rationality and truth. So we are left without an explanation of what Gert takes this relation to be. If I am right that the objective rational status of an action is tied to human nature, since it does not make sense to talk of objective rational status as divorced from the nature of the individual (or type of individual) for whom it is or is not rational (if not our nature, in virtue of what else could some action be rational or irrational for us?), then perhaps particularists have some options.

Obviously, if Gert's argument against drawing the analogy fails, then the analogy can still do work for particularists. However, as will be evident by the end of this dissertation, I do not rely on this analogy. Indeed, I ultimately reject the weighing model that is implicit in this analogy. There is, yet, much work to do to get us to that point. In the following chapter, I turn my attention to another candidate particularist view - that offered by David McNaughton and Piers Rawling in a series of papers.[17]

[17]See especially McNaughton and Rawling (2000, 2003, 2008, 2009)

CHAPTER 3

PARTICULARIST ALTERNATIVES I

3.1 Introduction

In Chapter 2, I attempted to show why Jonathan Dancy's particularism, while making progress toward establishing a particularist framework, fails in the face of objections from Selim Berker and Joshua Gert which attack his reliance on a radical holism in the theory of reasons. In short, Dancy's particularism allegedly fails because his commitment to such a radical holism leads to incoherence about practical reasons. One of the professed benefits of particularism over generalism is that particularists recognize the need for sensitivity to judgment regarding the particular reasons bearing on a particular case. The degree to which Dancy relies on holism in his attempt to accommodate this need seems to lead to trouble for his view. In this chapter, I will turn my attention to two other varieties of particularism in an effort to find a view more moderate than Dancy's in respect of its commitment to holism, but at the same time, less reliant on principles than full-blown generalism.[1]

To begin with, I make a few preliminary remarks about holism and its status in the pro-particularist literature. I then discuss McNaughton and Rawling's conception of holism. I explain how their holism has its roots in the holism offered by G.E. Moore (1903) in his *Principia Ethica*. Discussing this view, I think, helps to elucidate McNaughton and Rawling's holism, as it acts as a foil. In the final section of this chapter, I discuss a feature of McNaughton and Rawling's position with which I am in disagreement. They adopt what they call 'weak principles' - features which render their view unacceptably generalist. That is, insofar as one is looking for a particularist view and to the extent that McNaughton and Rawling are content having generalist features as integral parts of their view, the view is not a viable candidate.

[1] As previously mentioned, one of the views I will examine, as put forward by Margaret Olivia Little and Mark Norris Lance, is more accurately called a 'moral contextualist' view. Though, because their view holds its roots in particularism, it seems to me well worth considering in this context. Should particularism prove to be irredeemably flawed, presumably particularists would shift to a camp that is as near their own view as possible. So, if particularism is flawed, Lance and Little's moral contextualism may be a candidate for the next best thing.

3.2 A Brief Overview of McNaughton and Rawling's Particularism

One of the most jarring aspects of Dancy's particularism is the degree to which he privileges context-dependent features over invariant ones. That is, on Dancy's view, holism runs through reasons so thoroughly that even thick moral concepts such as courage, justice, beneficence, etc. can, in the appropriate circumstances change valence. If the circumstances are just right, the fact that it would be courageous to perform some act could, at least in principle, count against performing it.

This feature of Dancy's view is at the heart of some of the debates he has had to have over the years with, not just generalists, but others sympathetic to the particularist project as well. If the valence of even thick moral considerations can switch, then there seems to be no way to delineate those properties, considerations, or concepts that we tend to think of as particularly morally relevant from those that are either only sometimes, or never, morally relevant. Dancy must deny that there is anything special about the cross-contextual valence of the alleged thick moral concepts.[2] The thought is that if we have a basis for picking out certain properties as particularly morally relevant, then we have some basis, viz. the particularly moral content of those properties, for generating moral principles. On the other hand, if Dancy is correct about the holism of reasons and the extent to which reasons behave holistically, then, the idea is, we cannot articulate any moral principles describing the connection between natural properties or states of affairs and moral properties. That is, if Dancy's version of holism is correct, then whether or not the performance of some act has this or that moral valence is unpredictable, and so, inarticulable as a principle. This is the upshot of Dancy's particularism, then. Whether any consideration is going to favor or count against some action is not a question that can be answered from outside of a particular situation.[3]

This view that the moral valence of a consideration is so dependent upon the context in which it finds itself is a rather radical form of holism. Given Dancy's view, it is difficult to understand

[2]There is some discussion in the literature about whether it is concepts or properties that are thick. (See, for instance, Eklund 2011) I do not think that anything particularly important hinges on my selecting either concepts or properties as the relevant thick elements in the discussion. I may, therefore, switch back and forth between talking about thick concepts and thick properties.

[3]We have, of course, just discussed some problems with this sort of thoroughgoing holism in Chapter 2. I carry on in this chapter discussing Dancy's view the way I do, not because I think Gert (2007) and Berker (2007) have missed the mark. Rather, since Dancy is not the focus of this chapter, whether Gert and Berker are right is somewhat beside the point.

just how to make sense of the moral valence that a consideration itself has, since we only ever see considerations as parts of particular contexts and those contexts can influence the valence of the considerations in various ways. Dancy's view seems to be somewhat like what we might call a naïve view of the color of objects. We see a red ball and we might say, "That ball is red."[4] However, we only see the ball as red because of the way our visual apparatus functions in the lighting conditions under which we see the ball. If either our apparatus or the lighting change, then we might very well see the ball as having a different color. One problem for Dancy, then, might be like a problem for this naïve theory of the color of objects. We want to know the color of the ball itself, if indeed the ball itself has a color, but the color we see is so very dependent upon the context in which we see it. Similarly, we want to know what valence a particular consideration has; we want to know whether X gives us a reason to A or a reason not-to-A. The trouble is, we never come across X in isolation. We only ever come across it as embedded in a context - a context that can (and perhaps actually does) influence its valence.[5]

Holism comes in other, slightly less radical forms, though. McNaughton and Rawling are holists, but, unlike Dancy, they do not think holism commits them to the extreme context-dependence of all reasons. While Dancy discusses holism in great detail, outlining the different forms of relevance considerations might take, explaining the ways in which these different considerations can interact with one another to form reasons, etc. McNaughton and Rawling consider the matter in somewhat broader strokes and, occasionally in slightly different terms.[6] Given the probable connection between value and reasons - and given the probable tightness of this connection - it would be nice if it were relatively easy to translate the value-talk into reasons-talk without losing much, if anything, in the translation. McNaughton and Rawling say,

> It is plausible to suppose that value is a mark of reasons: if some entity or state is of positive value, then there is reason to take a positive stance toward it - to admire it, approve of it, bring it about etc. And, correspondingly, if it is of negative value, then

[4]Note, here we should hear our speaker as saying something about the 'true' color of the thing - the ball - rather than saying something about the way in which most humans in most situations *see* the color of the ball. In normal, casual contexts, if we want to determine whether someone is seeing a thing's color correctly, we can do this easily. What we do is we take a poll of the way people around us see the object in what we might call 'normal lighting conditions,' viz. outdoors, under a slightly overcast sky, around mid-day.

[5]I take this way of articulating the problem to be a simplified way of articulating the concerns raised by Gert (2007) and Berker (2007).

[6]See, for instance, McNaughton and Rawling (2000) and for their holism about *value* rather than reasons, see McNaughton and Rawling (2008).

> there is reason to take a negative stance toward it. Perhaps the link between value and reasons also runs in the reverse direction: if we have reason to take a positive stance toward something, then it is of positive value (mutatis mutandis for the negative case). (2008: 170)

Roughly, the idea, then, is that in acting we promote certain values by bringing about states of affairs and thereby generate, protect, demonstrate, etc. various values. Among other things, what we have reason to do is that which will promote value and what we have most reason to do is to promote that state of affairs that has the most value.[7]

As a brief aside, we might note that this sounds rather generalist in nature. Perhaps there is a maxim that looks something like, "Act always so as to bring about the most valuable state of affairs." However, while a simple, complete codification of (perhaps) morality, it is no more helpful than the claim that we ought to do what is right.[8] It offers no real guidance. When in the midst of certain moral conflicts, it is not always clear which option is right - and that is what we need help figuring out; this is surely one of the desiderata of moral principles. Moreover, even if one thought this toy maxim requiring we bring about the most valuable state of affairs were a correct and complete codification of morality, one would still have to offer some way of handling cases in which there were not a single course of action that produces the most value - say, because there is a tie between two or more courses of action. In such a case something other than value itself has to settle the matter in order for one's action not to have been arbitrarily performed. Otherwise, the choice situation is effectively a Buridan's Ass case.[9] McNaughton and Rawling need not worry about such cases, however, as on their view a reason's strength is determined by more than just value. So much for that aside - back to holism.

3.3 McNaughton and Rawling on Moore

McNaughton and Rawling tackle holism about value by taking G.E. Moore's famous discussion of intrinsic value and organic unities as a starting point. In Moore's discussion of intrinsic value in *Principia Ethica*, he supposes that if an object (or state of affairs) is truly intrinsically valuable,

[7] The value here promoted need not always be specifically *moral* value, as practical rationality may sometimes require we perform immoral actions.

[8] I say "perhaps" moral, here, because morality might require that we bring about the state of affairs with the most *moral* value, though this state of affairs may have less total value than an alternative.

[9] The assumption here, of course, is that there is nothing about the *kind* of value produced in either action that renders it a preferable choice for whatever reason.

then if we were to imagine it were the only thing in existence, we would still judge it to be valuable. Since Moore did think that we ought to perform that action that would promote the most (and best?) value, then in order to figure out which among the many courses of action before us would promote the most value we would have to add together all of the valuable elements of each course of action. In doing this we thereby determine the overall value of each course of action. Matters were slightly more complicated than this, though, for Moore. Moore believed that the value of a state of affairs is not necessarily equal to the value of the sum of its parts. He thought there was a distinction to be drawn between the value of a state of affairs *as* a whole, on the one hand, and the value of a state of affairs *on* the whole, on the other hand. This distinction is an important one. For Moore, the coming together of the elements of a state of affairs itself (the coming together) has a value. This is his value *as* a whole. On the other hand, the value of all of the elements of a state of affairs together with the value of those elements being together - the value *as* a whole - is the state of affairs' value *on* the whole. So, there are at least three important values here: the values of a state of affairs' parts, the value of the coming together of those parts, and the value of the parts plus the coming together of those parts. McNaughton and Rawling employ an example that illustrates this quite nicely.

Suppose Fred committed a crime and was caught and imprisoned by the authorities. We can assign values to, (i) his having committed a crime and (ii) his having been imprisoned by the authorities. However, we can also assign a value to, (iii) the state of affairs of justice having been meted out. That is, we might think that crime is bad, i.e. has negative value (so (i) is negative). Likewise, punishment, to the extent that it inflicts some pain on the criminal, is bad (value (ii)). However, the state of affairs in which just punishment is carried out against those who commit crimes is good (or at least better), i.e. has positive value (or at least less negative value) (this is (iii)). A plausible explanation of how we get such a state of affairs is that there is value in justice - which is distinct from the value of either the crime or the punishment. (McNaughton and Rawling 2008: 166-69) So we have:

> Value of the crime = -8
> Value of the punishment = -5
> Value of the combination (value as a whole) = +6
> Value on the whole = -7 (2008: 167)

It is important to note that the value as a whole is not a function of any of the individual values of the elements of the whole. Indeed, McNaughton and Rawling, in justifying their interpretation of Moore's position demonstrate how to understand some of his other claims. One such claim is that while (just) punishment following a crime always makes for a state of affairs that is better as a whole, nevertheless, depending on the crime and the punishment, the resulting state of affairs might be worse on the whole. So, taking the above numerical illustration, if the punishment inflicted were more severe, having a negative value of, they suggest, -7, this does not mean that the value as a whole would change. As McNaughton and Rawling have laid the example out, the value as a whole (might) remain the same: +6. The value on the whole in this case, though, is -9, whereas in the above example it is -7. (2008: 167-68)

There are, however, several problems with Moore's view. One problem, which McNaughton and Rawling discuss, is that there are some cases in which Moore's method of determining the value of some aspect of a state of affairs (a more circumscribed state of affairs) - the isolation test - does not make good sense. For instance, back to the above example of crime followed by punishment, McNaughton and Rawling wonder how we are to determine the value of the punishment *itself*. The punishment would be some harm inflicted, but without a prior crime. But then in virtue of what is it a punishment? For what would it be a punishment? That is, the value of the punishment *depends* on the existence of the crime for which it is a punishment. (2008: 168) Because the isolation test cannot help us to determine the value of a state of affairs, such as punishment, which depends for its value on another state of affairs from which it must be isolated, according to the isolation test itself, the test is defective.

A further problem with Moore's view is related to the connection that likely exists between value and reasons. A plausible view of this connection would maintain that, generally, if some object or state of affairs is valuable, we have (perhaps in proportion to its value) reasons to protect it, promote it, produce it, etc. If this is plausible, and in particular if it is plausible that the strength of our reasons vary with the degree of value that is up for protection, promotion, or production (perhaps giving special weighted consideration to certain dimensions of value), then we ought to be able to make sense of a view of reasons analogous to Moore's view of value.[10] In the rest of this

[10]I say "perhaps giving special weighted consideration to certain dimensions of value" because, depending on certain assumptions about the differences between certain varieties of consequentialism versus certain varieties of deontology, certain values may correspond to different types of reasons. For instance, it might be an open question whether

section, I examine what such a view would look like. As it will turn out, it is unclear whether such a view makes sense. To the extent that there is a tight connection between reasons and value it looks like Moore's view, and perhaps views taking Moore as their starting point, might be defective.

In an attempt to avoid any problems with punishment analogous to the one mentioned by McNaughton and Rawling's example above, I will look at a different example. Suppose I am deliberating whether or not to take a friend to the airport. That I promised to take him, that I owe him a favor, and that he has an important meeting later today at his destination all count in favor of my taking him to the airport. However, the fact that I have a lot of work that needs getting done and the fact that my car is in disrepair speaks against my taking him. We can, again, assign numerical values to these various considerations:

I promised him I would take him = +7
I owe him a favor = +4
He has an important meeting later today = +2
I have a lot of work to do today = -5
My car is on its last legs = -3

Weighing the case for some action and the case against that action against one another can be carried out in the following way: we sum the reasons for, we sum the reasons against, we weigh the case-for against the case-against, we then arrive at what we might call our all-things-considered-reason. Then we have an end to our deliberation - we know what we ought to do. So, summing the above values to determine the cases for and against, we have:

Case-for = +13
Case-against = -8

But here we come upon a wrinkle when we try to translate Moore's position into one in terms of reasons. In the case of value, the value of the combination of objects (or states of affairs) itself can, and needs to be, assessed. This is Moore's 'value as a whole'. Above, we talked about the value of a state of affairs of just punishment being meted out being a value in addition to the value of the offense and the value of the punishment for that offense. Similarly, in this case, we plausibly

value is to be construed monistically or pluralistically. The answer to this question might determine what kinds of reasons count in what way. I recognize that these are complicated issues with a large literature detailing all of the ins and outs of the debate. This project would be a very different one were I to take on these issues to even a small degree. So, I will set them to one side. I do not think much of what follows hangs on the details of these questions.

have to take into account, in addition to all the values already mentioned, the value of the state of affairs of a promise being kept and the value of the state of affairs or a promise being broken. The question is: What, in the case of reasons, is analogous to the value *as* a whole? What is the "case (for or against) as a whole"? Since, if we are to follow what seems to me a standard model of weighing reasons, once we have the case-for and the case-against, we have only to look at which is weightier to see what (given the available reasons) we should do. There is no more calculating, as it were, to be done. It is a simple matter of seeing which (admittedly artificial) number is larger.

What we have, then, is the case-for, the case-against, and each of these as-a-whole:

Case-for-as-a-whole = X
Case-against-as-a-whole = Y

Rather than make up arbitrary numbers, since value-as-a-whole seems not to have a clear connection to the values of the members of the whole, I will just use X and Y.

Now, in what could the case-for-as-a-whole actually consist? In the case of value, it at least seems to make some sense to talk about the value of the state of affairs of, say, just punishment for a crime being meted out. When talking about our reasons to act in such-and-such a way or not, though, there does not seem to be much sense to the idea of a 'case-for-as-a-whole' if it means anything more than the 'case-for'. Of course, when construing value as Moore does, it does mean something more than the sum of the values of which the whole consists. There seems to be a problem, then, with this way of drawing out the analogy. Perhaps this is a problem due to the analogical view itself - perhaps we cannot translate value-talk into reasons-talk in this way - or perhaps, instead, the problem is with *Moore's* view about value - that *it* cannot be translated into reasons-talk. Regardless of which it is, McNaughton and Rawling raise another concern with Moore's view of value that is analogous to a concern many have about certain ways of conceiving of reasons. So even if the preceding criticisms do not, by themselves, justify moving away from speaking of value and reasons as Moore does, there are still good reasons to do so.

Given Moore's reliance on the isolation test to assess the value of objects (or states of affairs), it seems that intrinsically valuable objects (or states of affairs) hold their value essentially. This is because in any situation in which we want to assess the value of some F, we must imagine F as the only entity in existence. No matter where, or what else is around (other properties, objects, states of affairs, etc.) when we assess F's value, we are always assessing the same thing - F alone

in existence. As McNaughton and Rawling say, from this we can arrive at a generalization: if here F is intrinsically valuable, then it always is, and to the same degree. (2008: 173) They claim that Moore's view results in both over- and under-generalization. It results in over-generalization because, despite what the isolation test seems to show, there seem to be certain objects (or states of affairs) that might not have the same intrinsic value in every whole of which they are a part. And it results in under-generalization because, like punishment, there seem to be certain things that cannot be isolated in the way required by the isolation test. As a result, Moore will be unable to generalize over these features and will therefore miss potentially important values.

3.4 McNaughton and Rawling on Holism's Scope

McNaughton and Rawling therefore reject Moore's holism, however not entirely. One important feature of Moore's view that McNaughton and Rawling believe he does get right is the notion that a whole can have a part that is necessary for that whole to have the value it does, while that part itself does not have its value necessarily.[11] (2008: 173) Nevertheless, unlike Moore, McNaughton and Rawling believe that the value of a whole is equal to the sum of the values of its parts. (2008: 174) This is because the value of a part, say, A, can alter the value of a different part, B, thereby altering the value of the whole, W, without A itself having its value necessarily. This is how their view is a holist one.

One of the reasons McNaughton and Rawling reject Moore's brand of holism is because it seems to lead to a certain sort of atomism (i.e. a variety of context-independence). Moore's view of value and his employment of the isolation test seem to entail certain objects (or states of affairs) have their value essentially, i.e. carry that value wherever they appear; they have context *in*sensitive value.

McNaughton and Rawling (2000)'s "Unprincipled Ethics" is in large part a discussion of the depth of holism. Whereas Dancy believes that holism is true of all reasons (even if there are certain reasons whose valence never actually varies),[12] McNaughton and Rawling maintain that there are

[11] This is a live option for Moore in part because he is committed to the fact that the value of a whole is not equal to sum of the values of its parts.

[12] Dancy holds both of these claims that admittedly seem, at first blush, to be in tension. He does so by claiming that there is a distinction between a reason and its content. (2004: 77) This distinction allows him to say that whereas a reason, by its nature, is holistic, its content, by its nature, never varies. But this invariance, again, is not due to the fact that we have a reason. Rather, it's due to the fact that we have a reason with such and such content.

certain reasons whose valence is context insensitive. That is, certain considerations, viz. some of those associated with the thick moral concepts, always count and count the same way whenever they appear. Note, though, that this is not to say that those considerations always count to the same *degree*. As we will see below, McNaughton and Rawling leave this option open.

The differences between McNaughton and Rawling's view and Dancy's are often subtle, and in the end these differences might not matter if particularism suffers from, as it were, endogenous defects. McNaughton and Rawling disagree with Dancy regarding both, what I am calling, holism's scope and its depth. Whereas Dancy thinks that it is possible for any consideration to undergo a complete valence switching (even if such valence switchings do not ever happen, his view is that they are possible), McNaughton and Rawling maintain that there is a set of considerations that, if holistic (and so possibly restricting holism's scope), they are only holistic in a limited sense (and so restricting the depth to which holism can affect these considerations). That is, this set of considerations, some of those associated with the thick moral concepts (often, the traditional virtues), are in some significant sense invariant. In this section I will examine what McNaughton and Rawling say about holism's scope and depth and how they claim their view is preferable to Dancy's. I believe that, if their view is significantly different from Dancy's (as, with respect to certain features they claim are different, the alleged difference may not be as significant as they claim), there are reasons to doubt whether these differences are good ones. That is, it seems that the ways in which they set out their view as distinct from Dancy's seem to make it a weaker version of particularism - and not simply in the sense that it is less extreme than Dancy's. Rather, it is weaker in the sense that it, either, is more vulnerable and therefore has to answer more questions before it is a viable alternative to Dancy's view, or, the view is in fact a generalist view in particularist clothing. Now, this latter concern may not be a *bad* thing. However, to the extent one has particularist intuitions, the more generalist a view is, the more conflict there will be between one's intuitions and the view in question.

Whereas Dancy believes that holism is true of all reasons (even if there are, as a contingent fact about our world, certain reasons whose valence never varies), McNaughton and Rawling maintain that there are certain reasons whose valence is, if not wholly context insensitive, then whose sensitivity to context is, as it were, rather dulled. Again, McNaughton and Rawling believe that at least some thick moral concepts are invariant - they never undergo a complete valence switching -

though their strength can be intensified or attenuated by context, they never change from counting for (against) an action to counting against (for) it. (2000: 273) So, for instance, the fact that some act is just always counts in its favor, though in certain situations the fact that an act is just only counts a little in its favor, while in others it counts very much in its favor. In the latter scenario, it would take a lot for it to be the case that I ought not to perform the just act, whereas in the former, since the fact that the act is just only counts a little in its favor, it would more easily be overridden. Again, though it is possible for the fact that the act is just to be overridden, it is never the case, according to McNaughton and Rawling and contra Dancy, that the fact that the act is just will count against performing it.

McNaughton and Rawling (2000) cast the debate between themselves and Dancy as one between competing forms of 'intuitionism', where 'intuitionism' does not refer to the sometimes spooky view that we have a moral-perceptual apparatus with which we 'intuit' moral facts from the world. Rather, McNaughton and Rawling are working with a different sense of 'intuitionism; viz. as "the view that there is an irreducible multiplicity of morally relevant considerations that have to be weighed to reach a moral verdict." (2000: 256) McNaughton and Rawling call their view 'thick intuitionism' to reflect the fact that they believe that "there are non-trivial cases of universally and counterfactually invariant valence, and these all involve thick moral properties."[13] (2000: 261) The variety of intuitionism that McNaughton and Rawling read Dancy as advocating, they call 'thin intuitionism', to reflect the fact that he believes that the only cases of universally and counterfactually invariant valence are cases involving the thin moral properties, right, good, wrong, bad, etc. (2000: 262)

The challenge that Dancy's thin intuitionism poses looks something like the following: it seems obvious that at least some non-moral considerations have variant valence. Additionally, the view that all such moral and non-moral considerations are of a piece is not entirely implausible. Indeed, one potential explanation of the difficulty in articulating the distinction between the moral and the non-moral is that moral reasons and non-moral reasons are different in subject-matter only, and not in a more substantial way having to do with their structure or logic. Because so many reasons are

[13]McNaughton and Rawling here refer to thick *properties* rather than, as it is usually done in the literature, to thick *concepts*. The distinction need not concern us here, however, it is worth noting that, as Eklund (2011) does, talking about thick *properties* renders it difficult to draw a distinction between thick and thin concepts. If the distinction is supposed to be one between items that are purely evaluative and those that are both evaluative and descriptive, then talk of thin *properties* would be difficult to pull off, as any thin property is arguably also tightly related to, if not identical with, a descriptive property. Again, we can pass over this without too much worry.

sensitive to context, what reason is there for thinking that any consideration has invariant valence across all possible cases? (2000: 263)

McNaughton and Rawling's response to this challenge is in part to remind us of their view regarding holism's depth. Unlike Dancy, they maintain that there are certain considerations - thick moral considerations - which behave holistically only to a limited extent. Take justice, for instance. On McNaughton and Rawling's view, that an act would be just always counts in its favor. However, McNaughton and Rawling recognize that context matters and that justice, like other considerations, is sensitive to context. They also believe, however, that justice is not as sensitive as some non-moral considerations. Other considerations in a given context can change the strength of considerations of justice, but these other considerations cannot ever change the just considerations so much that the just consideration counts against acting. So while the relevance of the fact that an act is just can change from context to context, the fact that the act is just will always count in favor of that act - it will never not count, nor will it ever count against performing the act. (2000: 273)

A further response to the thin intuitionist's challenge is to maintain that with respect to the variance of some thick concepts, the burden of proof lies with the thin intuitionist. She must give us a reason to believe that the valence of thick properties, such as justice, can vary. As McNaughton and Rawling say, "...it is a weakness of the position that its radical holism marks no distinction between the thick and the non-moral." (2000: 273) McNaughton and Rawling's thick intuitionism does not have this weakness. As they describe their view, thick concepts' invariant valence is what allows us to make sense of our ability to group acts together under the headings provided by those thick concepts. A further reason to prefer their view to the thin intuitionist view is that they have, what is for them, a satisfactory explanation of the centrality of certain properties - thick properties - that seems unavailable to the thin intuitionist. (2000: 266) For instance, the thick intuitionist can cite the cruelty of a certain set of stabbings - can point to the fact that this group of stabbings are all cruel - as an explanation of why they all matter morally. The thick intuitionist can then explain a stabbing that matters less, morally, for instance, one done in self-defense or the stabbing of a piece of fruit, by citing the fact it was not cruel. But the thin intuitionist cannot, because she has no good explanation of why certain non-moral properties matter, morally, more than others. Perhaps this requires a bit more of an explanation.

The main point of distinction between thin and thick intuitionists has to do with whether or not there are instances of universally and counterfactually invariant valence, and if so, whether and what there is to unify those invariant properties. The thick intuitionist claims that there is a set of such properties and they all involve thick moral properties: courage, justice, cruelty, dishonesty, etc. How exactly to characterize thick concepts is a notoriously tricky matter, but perhaps it is worth saying a few things. Pekka Väyrynen writes, "...thick terms and concepts somehow 'hold together' evaluation and non-evaluative description, whereas thin terms and concepts are somehow more purely evaluative..." (Väyrynen 2012: 235) Thick properties, then, are Janus-faced - they have an evaluative aspect and a non-evaluative aspect. Thick moral properties have a moral aspect and a non-moral aspect. Take justice, for instance. About justice, McNaughton and Rawling say, "Justice is a moral concept, and we suggest that understanding it, qua supervening term, requires the apprehension of its essential connection to the right." (2000: 266) If justice has an essential connection to the right, then it is clearly, at least *ceteris paribus*, more relevant, morally speaking, than, say, whether the Queen of England is wearing red shoes or blue shoes today. This essential connection is, presumably, due to the fact that as part of the concept justice, there just is an evaluative aspect - the fairness of the distribution, say, or the deservedness of the punishment - that gives the concept its moral tint. The idea is that there are configurations of non-moral (or non-evaluative) properties that are all related in some way, but perhaps not clearly at the non-moral level. What ties this set together is the fact that supervening on the non-moral (or non-evaluative) properties is the property: justice. And the fact that each member of this set of non-moral property configurations is a configuration of properties in which justice is at play is what allows us to group the property configurations into a set - even though at the purely descriptive level there may not appear to be anything that ties the property configurations together. This sort of thing is ubiquitous - we group things based on supervening or higher level properties, despite the fact that there seems to be nothing holding the items in the group together at the level supervened upon or at some lower level. For instance, at the physical level there seems to be nothing to keep us from putting absolutely everything in the same group - books, phones, cats, asteroids - they are all just so many atoms. But because there are these higher level properties, we can put these items into separate groups.

McNaughton and Rawling claim that thick intuitionism accords with our intuitions better than thin intuitionism because the thin intuitionist cannot explain why certain non-moral properties are more relevant in moral situations than others. Thick intuitionists can explain this because they maintain the invariance of certain thick properties. Naturally, one way the thin intuitionist might respond is by showing how the thick properties that McNaughton and Rawling claim are invariant are, in fact, not invariant. However, McNaughton and Rawling preempt such a response. They claim that any attempt to show that thick intuitionism is false by appeal to an example of a thick property with variant valence would only serve to show that the property in question was erroneously put on the list of invariant properties. It would not show that there are no such properties. (2000: 267) As they put it, "[t]he strategy of supplying examples is thus powerless in itself to show that there are no [invariant thick moral properties]." (ibid.)

But this response seems, in a way, just to point to the problem of induction. And despite the problem of induction, we nevertheless carry on relying on inductive reasoning, often with very good and reliable results. My point is merely that, yes, showing that just one or two of McNaughton and Rawling's invariantly valenced thick properties actually have variant valence may not be enough. However, repeated successes showing that some thick property's valence can shift, would shed a negative light on their view. Eventually, if the examples work, they will run out of thick concepts to pick out as having invariant valence.[14]

So then let's have a look at a few of the examples McNaughton and Rawling employ to examine whether or not there really are convincing examples of invariant thick properties. Promise-keeping, they note, looks as though we can spell it out in purely non-moral terms. For example, by saying a series of words, one makes a promise and thereby puts oneself under an obligation to fulfill that promise because of what those words mean. However, this is not the whole story. Because of the apparent fact that some promises are non-binding, in order to explain the normative force of promise-keeping, we need to do more than just say, "by saying a certain series of words" one puts oneself under an obligation to fulfill a promise. Otherwise our explanation of the normative force of promise-keeping is leaving something important out - viz. an account of when the conditions that normally hold for putting oneself under an obligation to fulfill a promise do not do so. So, for instance, when one makes a coerced promise or, McNaughton and Rawling think, when the promise

[14] Admittedly, this is a big 'if'. As is often the case in example-driven debates, the examples tend not to seem convincing to anyone who is not already convinced.

is to perform an immoral action, the promise has no real force - it is non-binding. 'Coercion' and, obviously, 'immoral' are evaluative notions, so, McNaughton and Rawling claim, thick intuitionism is best suited (or better suited than thin intuitionism) to spell out the role of promise-keeping in our moral lives. (2000: 269-70) Rather than a simple principle such as, "You ought to keep your promises," which picks out the invariant positive valence of promise-keeping, if we are going to have a principle of promise-keeping, we need something that accounts for the conditions under which it is not the case that you ought to keep your promises. So, this principle might look something more like, 'You ought to keep your promises unless the promise was extracted under coercion or is a promise to perform some immoral action'. It's important to keep in mind, of course, that this principle is not meant to be a so-called 'strong' principle - viz. a peremptory principle - rather it is a weak (i.e. contributory) principle. Strong, or peremptory, principles are verdictive. If a strong principle holds in a case, then no other principles can contradict it, no other considerations can bear weight against the course of action pre- or proscribed by the principle. It carries the day. Weak, or contributory, principles, on the other hand, apply in a more compromising fashion. Many weak principles can apply in a single case. The course of action required of us is that which is favored by the most or strongest principles. Weak principles can be overriden by other principles.

At this point in the discussion, the particularist presents the principlist with an alleged counter-example, after which the principlist, in an effort to defend the principle of promise-keeping might reformulate it to account for the exceptions highlighted by the particularist's example. The particularist usually follows this, though, with another clever counter-example, requiring yet another modification of the principle. This trend does not worry McNaughton and Rawling. They say,

> Even if our account of the [exceptions] is not complete and someone were to come up with an ingenious counter-example requiring supplementation or amendment to the principle, we can be confident that the needed amendments would be variations on the kind of theme we have already seen. (2000: 270)

Borrowing from Ross's list of prima facie duties, another example of an invariant thick property McNaughton and Rawling employ has to do with the duty of gratitude to someone who has helped you. (2000: 271) Whether you owe gratitude to someone who has helped you depends, at least in part, on the circumstances of the assistance you received. For instance, McNaughton and Rawling claim that if some person performs an immoral action which enables her to help you, then you owe

her no duty of gratitude. So to capture the normative force of the duty of gratitude, we need to appeal to moral notions having to do with the circumstances within which the assistance for which you are grateful is offered. Because of the nature of the duty of gratitude, any excusing conditions are going to have to have, as a feature, some normative property that acts as a disabler of the way the duty of gratitude normally works.

The key here is that McNaughton and Rawling maintain that part of what it is to have a correct grasp of certain thick concepts is an understanding of those conditions when the properties associated with those concepts work normally and when they do not. A hard-line particularist, like Dancy (or, like Dancy used to be)[15] wants to maintain that the set of circumstances under which thick properties do not work the way they normally do are too numerous for human minds as they are to usefully grasp - hence their particularism. McNaughton and Rawling, however, maintain that "the list of [excusing] considerations is not open-ended, and it is knowable in advance. It has a rationale. The conditions are inherent in a proper understanding of the nature and role of [for instance] promising in the fabric of our moral life." (2000: 270) As Dancy notes at the beginning of *Ethics Without Principles*, one way to examine whether the list of excusing conditions really is open ended would be of offering counterexamples to proposed principles. The principlist then responds by offering amendments, to which more counterexamples are then offered, and so on until either we have a principle without counterexamples or until the principlist gives up. This might be a fun way to spend an afternoon, but it does not seem at all like the best way to go about the debate. So I will not pursue that strategy here, rather I will set that question to one side. If looking at particular examples of the variance or invariance of thick properties is not going to help settle the debate, what other reasons for preferring thick intuitionism are there?

One reason for preferring thick intuitionism is that, despite the thin intuitionist's claim that her view's thoroughgoing holism accords best with our moral intuitions, thick intuitionism itself is not inflexible. McNaughton and Rawling claim that, despite the fact that the valence of a thick concept cannot go from positive to negative, or from negative to positive, there is room on their view for a wide range of variance. For instance, while even when the relevant conditions are met, promise-breaking always has a negative valence, just how strongly promise-breaking counts against performing an action depends on a variety of circumstances. Those circumstances include, but

[15]Dancy has, in recent years, admitted that there are what he calls 'default reasons' - reasons which standardly operate with a particular valence.

are not limited to: to whom the promise was made, the content and context of the promise, the circumstances under which the promise was broken, etc. When someone ought to break a promise and does not, she does the wrong thing. But McNaughton and Rawling maintain that such a case is not one in which promise-breaking's valence changes from negative to positive. Rather, this is just a case in which another consideration outweighs promise-breaking's magnitude. (2000: 272) So, thick intuitionism provides us with the machinery to account for *apparent* valence-switchings without actually committing to the thin intuitionist's view.

Another reason McNaughton and Rawling claim we ought to prefer thick intuitionism over the thin variety is that the latter fails to do justice to the prominent role that thick moral concepts play in our lives. As discussed above, McNaughton and Rawling believe that with regard to certain thick concepts, the burden of proof lies with the thin intuitionist to show that their valences can vary. Ultimately, whether one sides with the thin or thick intuitionist here might depend on, among other things, whether one has a theoretical ax to grind or on who thumps the table most loudly. I must admit, determining whether one consideration merely overrides another or whether the former causes the latter's valence to switch can be rather tricky. Generally, parties to these sorts of debates take reasons to be facts and facts to be some sort of abstracta. It is, therefore, just about impossible to settle the matter by looking at the world. For that reason, it is unclear to me in virtue of what, other than the way matters strike us, we are supposed to decide which is more likely to be true: that some consideration's valence was switched or that it was overriden by some other consideration. For that reason, I think on this matter, neither thick nor thin intuitionism has an edge on the other.

Yet a third reason to prefer thick intuitionism to thin is that thick intuitionism gives us the best account of moral modality. Dancy's thin intuitionism attempts to account for the apparent centrality of certain properties and the apparent necessary relevance of those properties by appealing to the inference from $\Diamond$-p to $\Box\Diamond$-p in S5. On his view, we get necessity in morality and explain the role principles seem to play by recognizing that when we observe the role a property plays in a particular situation, we are observing a role that it is possible for that property to play. Given the characteristic axiom of S5, we then can infer that it is necessarily possible for that property to play that role. But this is just getting necessity on the cheap. McNaughton and Rawling counter by noting that this is "too easily paralleled in low-level non-moral cases." (2000: 274) Dancy is trying

to give an account of moral modality that ought to be acceptable to principlists who are wary of particularism's rejection of principles. But he fails to do this because of this very weak necessity.

On the other hand, McNaughton and Rawling's thick intuitionism provides a much more robust account of moral modality. Their weak principles, employing thick properties as they do, are expressions of the relation between thick and thin properties. This tight relation between (certain) thick and thin properties - it being part of what it is for something to be a thick concept or property - manifests itself in the invariance of those thick properties. And this invariance is just the sort of modality that principlists are looking for - it is built right into McNaughton and Rawling's thick intuitionism.

3.5 Weak Principles and Default Reasons

McNaughton and Rawling develop a view that's meant to be a particularist alternative to Dancy's particularism. Again, their view admits, contrary to Dancy's, of a series of 'weak moral principles' - contributory principles that pick out certain properties, those corresponding roughly with the thick moral concepts, as having invariant valence. And, again, there are at least three benefits, they claim, that their view affords which Dancy's does not: (i) Thick intuitionism is more flexible than a principlist view that employs strong principles because of the former's sensitivity to context, but not so flexible that it seems not to have any kind of structure, like Dancy's view, (ii) Thick intuitionism respects the important role that thick concepts play in our moral lives. Dancy's view does not do this, and (iii) Thick intuitionism provides a robust account of moral modality. Again, Dancy's view does not do this. So there seem to be a few good reasons for preferring McNaughton and Rawling's view to Dancy's view. However, there are also a few aspects of their view which might be cause for concern. In this section, I will discuss some of them. In short, some of the ways in which McNaughton and Rawling's thick intuitionism is distinguished from Dancy's thin intuitionism make the McNaughton and Rawling variant less and less distinctly particularist and more and more apparently generalist. These virtues of their view, then, may in fact be vices, if what we are after is a distinctly particularist moral theory. It is, however, worth mentioning a few modifications Dancy has made in recent works in which he makes some concessions to McNaughton and Rawling and other dissenting voices - concessions that bring their views more in alignment with

one another. Because of these modifications, some of the criticisms I level against McNaughton and Rawling should apply to Dancy as well. First, though, a brief aside.

3.5.1 Why 'More' Particularist?

What do I mean by 'distinctly particularist' in the previous paragraph? What justifies the push toward a view that is 'more particularist'? These are difficult questions. Here, and at other places throughout this dissertation I raise this charge against the views I am discussing. I therefore ought to make a few remarks as to what I mean and why I believe I am justified in seeking a 'strong particularism.'

As I have mentioned in various places throughout the preceding two chapters, there are a variety of particularist views available for adoption. Some are 'less radical' others are 'more radical' with regard to their stance on the role principles play, if there are any. That is, some are views according to which there are no true moral principles at all. On the other hand, there are views, like Dancy's, according to which whether or not there are principles (as part of the fabric of reality - the metaphysical makeup of the world) we do not need to *rely* on them to make moral distinctions or arrive at moral judgments. Now, really, whether we call such views more or less radical is a matter of our interests. Which views we take as more or less radical will depend on in what parts of moral reality or moral reasoning we are most interested. For my purposes, given that it is reasons, particularly contributory reasons, and how they weigh up together (i.e. within the GWF), I organize views according to the cross-contextual sensitivity of the reasons they espouse. So, if Berker (2007)'s reading of Dancy is correct - and I believe it is - then Dancy's view is rather radical. The reasons holism Dancy employs in his metanormative framework permits context to determine nearly everything relevant to the strength of all normative considerations.[16] McNaughton and Rawling, as I discuss in the present chapter, are less 'radical' than Dancy. Their view, permitting weak principles, accommodates invariant considerations. While their strength may vary, they never go from in favor of Φ-ing in one context to against Φ-ing in another context. Now, employing a variety of principle, as McNaughton and Rawling do, their view seems clearly less 'distinctly particularist' than Dancy's. At least, that is what I am calling it. The term 'particularism' was once used only to describe views which countenanced a shapelessness between the first two of three levels Berker

[16] Dancy permits some cross-contextual insensitivity with regard to his default reasons. These, however, have problems all of their own, so I will not discuss them further here. See McKeever and Ridge (2006) for a criticism of Dancy's default reasons

(2007) describes in Dancy's metanormative framework - the non-normative and contributory levels. McNaughton and Rawling, being in agreement with shapelessness at *that* level, clearly belong in the particularist camp. However, in modern terms, particularism is a term for views which eschew any codifiability of the moral domain. Because McNaughton and Rawling employ weak principles, then, their view is, again, *less particularist* than Dancy's.

Is there a reason to be *more* particularist than McNaughton and Rawling? I think the short answer is 'yes.' Given my own particularism and my conviction that generalism is flawed, I think that to the extent a view has generalist elements, it is a view that could be improved upon. The long(er) answer, which I will discuss in more detail in Chapter 5, is that I believe there is a fundamental problem with contributory reasons, given the GWF, which Berker (2007) and Gert (2007) have pointed to. To the extent a view employs contributory reasons and adheres to the GWF, it is subject to the concerns raised by both Gert (2007) and Berker (2007). All extant particularist views are such views. My aim is to find a view that is not subject to those concerns. The only way I can see to do that is to be 'more particularist' - to be more extreme. The trend in the literature, as I see in the shifts from Dancy (1993) to Dancy (2004) (especially the explicit embracing of, and the treatment given to, default reasons), has been to quiet generalist critiques of particularism by making concessions to the generalist - by making particularism more and more like generalism. That strategy has so far failed to provide us with an adequate particularism. This is why I am looking for a 'strong' particularism.

McNaughton and Rawling's main point of departure from Dancy's view was in their advocacy of weak moral principles and his rejection of any sort of principle. Without such an appeal, the particularist has a 'flattened' normative landscape - all normative considerations are on a par with one another, none more relevant or important than any others are.[17] This is, of course, in stark contrast to the way the normative domain, and perhaps morality in particular, seems to work. It does seem as though certain considerations count more and are more important, morally, than others. For instance, whether this particular speech act is the telling of a lie seems, *ceteris paribus*, as though it ought to matter more than whether my shirt is black or blue today. Similarly, in other normative domains, say, that of rationality, certain considerations seem to matter more than

[17]This 'problem' receives articulation in every admission, by particularists, that humdrum features such as shoelace color can matter, morally, as much as some thick property, like courage, to what we ought to do. An early articulation of this feature of particularism (or, more precisely, a feature of holism) as a problem for particularism appears in McKeever and Ridge (2006).

others do. Whether what I am about to drink is poisonous or not (and, assuming I am aware of this fact) seems to matter more than whether the shirt I am wearing is black or blue. According to the strong particularist, or, more appropriately perhaps, according to the extreme holist, any consideration's valence can change given the appropriate context, so no consideration, necessarily, carries its valence with it across context. By appealing to the notion of a default reason, Dancy hopes to alleviate concerns over the flattening of the normative landscape while still maintaining a rather strong holism. Default reasons are, roughly, considerations that carry their weight across context, but not necessarily. For instance, considerations of justice are default considerations for Dancy. That some action is just is usually a reason to perform that action, though not necessarily a decisive reason. However, whereas McNaughton and Rawling are inclined to say that this is always, and never not, the case, Dancy maintains that justice's status as a default does not make it immune to the effects of certain contexts. Dancy claims that, if it is the case that there are invariant reasons, they are invariant because of the reasons that they are - because of their content - and not because of their status as reasons. But notice that what this seems to imply is that reasons are these peculiar ontological objects, floating around contentless, until an appropriate state of affairs obtains to provide content to attach itself to the reason. Again, Dancy seems to have turned reasons into something like Locke's substratum - something I know not what, onto which attach normatively charged contents.[18]

Whether the distinction between a reason and its content is a genuine one or not, it is clear that Dancy wants to be able to countenance default reasons. Presumably, if it were to turn out that Dancy had committed himself (accidentally?) to a strange reasons-substratum view, he would

[18]McKeever and Ridge (2006) raise another worry for Dancy's default reasons. Dancy describes default reasons as those considerations which come 'switched on, as it were', in need of explanation for the way they function only when they function contrary to their default - only when there is a disabling consideration present that we can cite in an explanation of why the default did not here count the way it normally does. McKeever and Ridge point out, however, that this way of explaining the distinction between default reasons and non-default reasons relies on the apparent distinction between the presence of a disabler and the absence of an enabler. But this latter distinction is untenable. By way of example, McKeever and Ridge ask us to suppose the fact that an action will promote pleasure is a default reason - usually it is a reason but sometimes it is not. One of the reasons why the promotion of pleasure might not be a reason is that it might be sadistic, and this seems to be the sort of thing that would change the valence of a reason significantly. On any occasion on which the pleasure promoted is sadistic, we can cite the sadistic nature of the pleasure as a disabler to explain why the pleasure does not function as it normally does. However, there is another explanation available. We might instead say that the relevant enabler is the non-sadistic nature of pleasure - in which case when a certain instance of pleasure is sadistic, what we have is the absence of an enabling condition. If we opt for this characterization of the situation, then we must admit that the promotion of pleasure does not have the status of a default reason. But there seems to be no good reason to cite the presence of a disabling condition rather than the absence of an enabling condition given Dancy's view. See Dancy (2007) for Dancy's attempt at a reply to these concerns about default reasons.

alter the view in some way, or explain the problem away, to hold onto default reasons. Similarly, as McNaughton and Rawling (2000) discuss, they can explain away alleged counter-examples to their proposed weak principles by appealing to the sorts of responses they have already made in anticipation of other such objections. They write,

> Even if our account of the [exceptions to weak principles] is not complete and someone were to come up with an ingenious counter-example requiring supplementation or amendment to the principle, we can be confident that the needed amendments would be variations on the kind of theme we have already seen. (2000: 270)

I must admit, however, that I am unsure what justifies this confidence if not a prior, perhaps question-begging, commitment to the kind of invariability that is at issue. But suppose they are right, and any needed amendments to a given weak principle are of the same type as those amendments to the weak principle about promise-keeping. Also, suppose McNaughton and Rawling are right that, "...the list of [excusing] conditions is not open-ended, and it is knowable in advance...The conditions are inherent in a proper understanding of the nature and role of promising in the fabric of our moral life." (2000: 269-70) The question then seems to me, "How is this different from a moderate generalist theory?" We might forgive someone for thinking that a view that advocates principles and has ready to hand a series of remarks about exceptions to those principles, as well as advocating for a restricted holism, might be a generalist view. For instance, consider the fact that we can characterize even a minimally sophisticated utilitarian view as a view that advocates principles, has a story about exceptions to those principles, and advocates for a restricted holism. McKeever and Ridge (2006) briefly describe something like such a utilitarian view in their effort to refute Dancy's claim that the truth of holism gives us reason to believe that particularism is (at least likely) true. (McKeever and Ridge 2006: 29)[19]

McNaughton and Rawling's view is a distinctly particularist one in large part because their set of weak principles is not meant to tell the whole story. Even having a complete understanding of the weak principles for which they advocate, they claim, is "insufficient for the correct moral appreciation of the particular case." (2000: 256) In addition to an understanding of weak principles,

[19] In conversation and correspondence, David McNaughton has expressed the view that the list of exceptions to McKeever and Ridge's example of a simple holistic Utilitarianism is *too* restricted. That is, just the one exception would not qualify the view as properly holist. However, if we understand the view McKeever and Ridge outline as one according to which all reasons are ultimately derivative on considerations of pleasure and pain, then McNaughton's worry seems avoidable.

the correct apprehension or appreciation of a moral situation requires judgment. The role of judgment in such cases is to bring us from our apprehension of weak principles and the case at hand to a verdict on the matter. One might fully understand a weak principle that articulates the connection between cruelty and wrongness and yet, due to defective judgment, be unable to identify when an action, whether one's own or another's, is cruel. Understanding the weak principle does not entail one's being able to appreciate the case at hand in all of its complexity. So, unlike the toy-utilitarian view McKeever and Ridge employ in their argument against Dancy, McNaughton and Rawling still maintain that, given a full appreciation of the situation, one cannot just read off from their weak principles the correct moral verdict. Perhaps, then, the question is whether this 'ineliminable role for judgment' is sufficient to render a view a particularist one. After all, this seems to be the most significant theoretical apparatus that distinguishes McNaughton and Rawling from generalists. And yet, the fact that they embrace weak principles seems also to connect them to generalists in a way that might be uncomfortable for a particularist of a more anti-principlist stripe.

Some state of affairs S is composed of various features (facts, considerations), some of which are morally relevant, others non-moral. Suppose one aspect of S is my deliberating about whether to slow my car down for a group of undergraduates who are drunkenly stumbling across a street near campus. Ross would say, among other things, that because I have a *prima facie* duty of non-maleficence, all other things being equal, I ought to slow my car down. Similarly, McNaughton and Rawling would presumably say something like there is a weak principle, which, however it is worded, would require that I not run over inebriated undergraduates. But putting things this way makes the view look very much like a generalist view embracing hedged principles - for instance Pekka Väyrynen's view, which will be the subject of Chapter Four. This similarity seems particularly salient when one recalls that, in their discussion of promise-keeping, McNaughton and Rawling claim that an understanding of the conditions when a promise need or ought not to be kept is "inherent in a proper understanding of the nature and role of promising in the fabric of our moral life." (2000: 269-70) This seems to me like an admission that one cannot get on morally, when promises are involved, without being in possession of (having an understanding of) the moral principle governing promise-keeping. But this gives a role to this principle much greater than that

of a signpost – it makes it *necessary* for moral life. And that seems a role that ought not to exist within a particularist framework.

It is, of course, open to McNaughton and Rawling to suggest that their view is particularist in the sense that they do not believe that there can be non-normative characterizations of principles capturing all moral action-types. That is, there can be no such principles capturing all and only cruel actions. The features that make an action cruel are too many and varied to be capturable in a finite principle. They permit weak principles because all weak principles are already normative. Their view is, as it were, about the move from the non-normative level to, what Berker refers to as, the contributory level. McNaughton and Rawling are perfectly happy with the codification of morality from the contributory level to the overall level, i.e. with principles that tell us, given certain normative considerations, what we ought to do all-things-considered. However, it looks to me like McNaughton and Rawling might be begging an important question against the pure particularist when they say, as quoted above,

> Even if our account of the [exceptions to weak principles] is not complete and someone were to come up with an ingenious counter-example requiring supplementation or amendment to the principle, we can be confident that the needed amendments would be variations on the kind of theme we have already seen. (2000: 270)

Again, it is unclear what can justify this confidence.

In the next chapter, I suppose that, having found no joy looking at Dancy or McNaughton and Rawling's views, maybe I ought to give up on such a pure particularism. In that case, *perhaps* McNaughton and Rawling's view *would be* a viable candidate once more (if not for the question-begging just mentioned). Until that matter is settled, however, I suggest that McNaughton and Rawling's view is still not, in fact, a viable candidate. In Chapter Four, then, I will look at two Particularist-ish alternatives - the moral contextualism advocated by Margaret Olivia Little and Mark Norris Lance, and the moderate generalism advocated by Pekka Väyrynen.

CHAPTER 4

PARTICULARIST ALTERNATIVES II

4.1 Introduction

In Chapters 2 and 3, I tried to show how the two most prominent varieties of moral particularism are meant to work and how neither one is without fairly significant problems. In this chapter I turn my attention to two other, related views. If, contrary to (for instance) the anti-transcendental particularist's claims, moral thought and judgment do depend upon a suitable supply of moral principles, particularists are likely to prefer a variety of generalism that makes as few moves away from particularism as possible. As the next chapter will be a speculative examination of a different path to particularism (which, admittedly, may not ultimately succeed), in this chapter I look at the prospects for two moderate generalisms. As will become evident, both of these views have something to offer particularists should particularism's problems remain unresolved. Indeed, both views take the holism of reasons as one of their starting assumptions - just as particularists so frequently do.

First, in section 4.2, I will discuss the quasi-particularist view developed by Mark Norris Lance and Margaret Olivia Little. I call the view quasi-particularist because, while Little (2000) is expressly particularist, Lance and Little (2006a; 2006b; 2007, 2008) insist the term no longer accurately describes their view. The moral contextualist view they offer is like particularism in that it rejects exceptionless principles. However, unlike certain strains of particularism (such as Dancy's anti-transcendental particularism), it accepts and employs a role for a type of defeasible generalization.

In section 4.3, I will discuss a view that is also an expressly anti-particularist, generalist view. Pekka Väyrynen has recently developed a view exploiting what he sees as a natural assumption about the normative. Reflection on this natural assumption leads Väyrynen to the development of a sort of hedged principle. Both the views offered by Lance and Little and Väyrynen (2006, 2008, 2009, 2012) share some important features. Obviously, as they both accept a type of generalization,

they are similarly anti-particularistic. However, there are important differences between the views, discussion of which might be helpful to future developments of particularism.

One way to deal with objections to particularism might be to give up some ground to the generalist - to make some adjustments that make particularism more like a moderate generalism. This is how I view what Lance and Little have done in their shift from particularism to contextualism. If, however, the varieties of particularism and generalism exist on a spectrum we might worry that somewhere in the middle, the lines are too blurry for us to distinguish between particularism and generalism. Lance and Little are taking us in this direction. Similarly, Väyrynen's moderate generalism is at least as similar to particularism as it is to a more stringent variety of generalism, e.g. classical utilitarianism and deontology. In this chapter I will argue that the moves made to the center, towards moderation, are unsuccessful. In the next chapter I will argue that we ought not to be moving toward a middle-ground between particularism and generalism but instead, if we are going to find a variety of particularism that works, we may have to move back to the fringes.

4.2 Lance and Little

4.2.1 Preliminaries

Lance and Little develop a moral contextualist view that many particularists might find attractive should particularism prove to be, in principle, problematic. While Little (2000) advocated for a variety of particularism, Lance and Little have since decided that the term 'moral particularism' does not quite fit the view they are offering. Nevertheless, the view is, in some respects anti-generalist in the way particularism is. It is therefore an avenue for theorists who agree with the heart of the particularist claim about generalism, but who, perhaps, accept that particularism's problems are insurmountable. In this section, I will examine the view offered by Lance and Little in their series of papers. The view that Lance and Little end up with is a view that seems to be committed to certain claims that Little (2000) would have rejected. This by itself may not be a problem. But if Little (2000) rejected the claims for good reason, then the fact that Little and Lance's view accepts the claims seems a mark against it.

4.2.2 The View

Perhaps the best way to go about this discussion is to begin with Little's discussion of the view qua particularist view and then to describe the moves Lance and Little make and the reasons they cite for their rejection of particularism and preference for moral contextualism.

In Little (2000), Margaret Olivia Little gives the most comprehensive description of the reasons she is a particularist rather than a generalist. Frequently in Lance and Little's work, they start from the assumption that some form of particularism can be gotten to work, and they proceed to discuss their preferred form of particularism (or later, their preferred form of anti-exceptionless-generalism). Little argues, as other particularists have, that the motivating intuition behind particularism is reasons holism. Well, as these things tend to go, it is a little more complicated than that. Particularists believe in the shapelessness[1] of the moral with respect to the nonmoral. That is, as McNaughton and Rawling (2003) note with respect to tin-openers, the ways in which features of situations can combine to comprise, say, a cruel act cannot be coherently described without appealing to the notion of cruelty. In the case of tin-openers, for instance, we have metal ones with thin handles and metal ones with thick handles, we have ones with gears and ones with levers, we have all sorts, and the only thing that unifies all of the things that we call tin (or can) openers is the fact that they are tin openers. This is the lowest level of categorization to which we can appeal to unify the group. Similarly, with all of the acts that we call cruel, some are speech acts, some are overt bodily acts, some are actions of omission and we cannot unify the group of acts except by appealing to the notion of cruelty. This is in large part what it is for cruelty to be shapeless - there is not a recognizable shape, apart from and 'below' the level of being cruel, by which we can categorize cruel acts. There is no pattern to the configuration of wholly non-normative properties that, somehow, result in cruel acts. If this is right, then the prospects for the codification of morality seem to Little (2000) to be rather bleak.

How could we, for instance, articulate any principles capturing all the ways various actions can be cruel when such articulations would be so long that they would effectively be supervenience functions? Particularists do not deny the supervenience of the moral on the nonmoral. But this does not mean that we can get the sorts of principles we are looking for by pointing to a long disjunction containing all properties that can be in the subvenient base for a given moral property. Such

[1]Little attributes the term 'shapelessness' to Simon Blackburn. See Blackburn (1981).

supervenience functions are not useful or realistically informative, because they are too unwieldy. (Little 2000: 285)

Little claims that the truth of the holism of reasons means that "we simply shouldn't expect to find rules that mark out in nonmoral terms the sufficiency conditions for applying moral concepts." (2000: 284) She claims that the moral import a natural property carries depends, in a way that cannot be specified beforehand, on the background context in which we find that natural property. Part of the reason why we cannot specify the import beforehand is presumably because we are not omniscient creatures. But also, there are certain moral features of actions that simply cannot be spelled out without appealing to moral, or otherwise normative, notions. For instance, Little cites the example of rape. She rightly points out the difficulty one would have in giving a definition, in purely non-normative terms, of that which we condemn when we condemn a rape. She says,

> It is actually very difficult to cash out that kind of context — to define the act we mean to condemn — without helping ourselves to concepts that turn out, on reflection, to be irreducibly laden with the evaluative (think of trying to isolate the instantiation conditions of *consent* without the help of moral notions like *fair*, or of *force* at the level of physical mechanics). (2000: 284)

However, it is unclear to some whether holism in fact entails particularism. McKeever and Ridge argue that holism, as it is discussed in the literature, either supports particularism only trivially, because the two views are not properly distinguished, or else it provides particularism with no special support. (2006: 26ff.) Mark Schroeder (2011) argues that, rather than just not supporting particularism, "...normative claims are holistic because they are general, rather than because they defy generalization." (2011: 328) He claims that statements about reasons are pragmatically associated with facts about the weights of reasons, and these weights are general facts. Such statements are pragmatically associated with facts about the weights of reasons because, as Schroeder attempts to show with a series of examples, when we claim that there is a reason to do something, we generally mean that there is a *relatively weighty* reason to do that thing. Such facts of relative weight are *general* facts. (2011: 332) Schroeder claims that general facts are holistic. For instance, 'all swans are white' is a general fact. It is holistic in the sense that a partial specification of how the world is will not guarantee its truth - the truth of the claim can change given the presence or absence of some feature of the world. (2011: 330) So, if reasons claims are holistic, it is because they are general claims. However, Schroeder's point depends on the pragmatic association

of reason-claims with facts about the weights of reasons. Even if it is true that as a matter of fact, in practice our claims about reasons are claims about relatively weighty reasons, it does not follow that facts about reasons are facts about relatively weighty reasons. Schroeder is in fact very careful not to commit himself to this claim. (2011: 332) But that is the claim that he would need to make in order to target the varieties of particularism that I have been discussing.

One might expect that particularists themselves would be a bit more careful about properly describing holism and particularism and keeping the two distinct. However, Little (2000) does *not* seem to distinguish carefully between the two. So, as McKeever and Ridge have noted, Little gets holism's support for particularism in a trivial fashion. Under such an interpretation, holism just *is* particularism. Little says of the moral contribution made by nonmoral features of a situation that it is "holistically determined: it is itself dependent, *in a way that escapes useful or finite articulation*, on what other nonmoral features are present or absent." (2000: 280, my emphasis) But particularism is meant to be a claim about the connections between nonmoral features of a situation and the moral features to which they give rise. If that is what holism is, then what is there left for particularism to be? Whether nonmoral properties behave holistically or not - whether they are heavily context-dependent - is an entirely separate question from whether or not that context-dependence can be codified. The fact that Little so seamlessly and effortlessly moves from talking about holism to talking about particularism seems to me a regrettable black mark against the particularist camp. Little suggests that holism (and, I suppose, therefore particularism) is obviously true and widely accepted as true in other normative domains. If that is the case, and if there are no good reasons to suspect that morality is fundamentally different from those domains, then this might lend support to the particularist position (at least, if those other domains happen also not to be codifiable in the way the particularist claims morality is).[2]

Little points to aesthetics and to epistemic justification as evidence that holism is accepted in other domains. For instance, Little notes that "the bold stroke of red that helps balance one painting would be the ruin of another" and there is no way to specify in non-aesthetic terms the conditions in which it will help and the conditions in which it will detract." (2000: 280)[3] Similarly, in the case of epistemic justification, Little says, "beliefs and experiences do not carry

[2]Gert (2008) argues that there are important and relevant differences between the epistemic domain and the moral domain which render the analogy between the two illegitimate.

[3]Little credits a conversation with David McNaughton for this aesthetic example.

their justificatory import atomistically," whether the experience one has of seeing a table counts as evidence that there is a table there depends on the context of the experience - for instance, whether one has taken psychotropic drugs recently. "There is no way to codify the conditions under which an experience as of seeing a table is evidence for there being a table." (2000: 281)[4] These two examples seem like familiar, or at least not terribly controversial, examples of non-moral domains in which holism is present. And so the argument is, if holism is true in these other domains and there are no substantive differences between these other domains and morality, then we ought to think that holism is true about moral reasons too - we ought to think that it is possible for a consideration that acts as a moral reason for Φ-ing in one situation to undergo a switch in valence. But then, again, the truth of holism may not entail the truth of particularism. Perhaps even further, if Mark Schroeder (2011) is correct and "...normative claims are holistic because they are general, rather than because they defy generalization," then holism is evidence *against* particularism. However, as I hope to demonstrate in the next chapter, whether holism is true or not may be irrelevant to the truth of particularism, as long as we are willing to revise our conception of what practical reasons are. So, for now, I shall set the issue of whether holism actually supports particularism, and if so, how, to one side.

And now, a gentle metaphilosophical aside: it is very difficult to build an argument from a place free from assumptions. But generally, I take it, we try to start from territory that is as neutral as we can get. By doing this, I take it, we are trying to give as little reason as possible to those with views very different from our own to reject our arguments for reasons having to do with our starting point. Unfortunately, Little makes a fair few contentious claims. She begins by claiming that "the model backing particularism clearly belongs in the non-naturalist camp." (2000: 279) She says this because particularists claim that the shapelessness thesis is at the heart of their particularism. But to claim that if one thinks the shapelessness thesis is true, then one is a non-naturalist rules out the possibility of one being, for instance, a non-reductive naturalist according to which, moral properties are natural properties, but nevertheless not reducible to natural, non-moral properties. So, for instance, cruelty, might itself be a natural property. Such a theorist could also be a particularist because of the shapelessness of the moral with respect to the nonmoral. Shapelessness need not entail non-naturalism as long as non-reductive naturalism is permitted.

[4]Note again that Little is sliding seamlessly from talk of the context-dependence of some consideration to talk of the uncodifiability of the domain governing that consideration.

It is difficult to discern a clear argument in Little (2000) for particularism (or contextualism). Rather, Little tells us why she is an anti-generalist in the way she is, providing examples and counterexamples, and she tells us something about what follows if one is an anti-generalist. However, because of the significant shifts in her view in subsequent work co-authored by Mark Lance, it would perhaps serve us best to look to the latter for a description of the consequences of their brand of anti-generalism.

4.2.3 Defeasible Generalizations

In a series of papers, Lance and Little, have laid out their positive particularist (and later contextualist) theory according to which there are no finite, useful, exceptionless generalizations, but there are finite, useful, exception-ridden generalizations. Indeed, as they put it, "disciplines from epistemology to biology, from ethics to semantics, are rife with generalizations that seem explanatory even while they are porous - shot through with exceptions that cannot be usefully eliminated." (Lance and Little 2008: 61) All of the following generalizations which Lance and Little mention are on a par, despite being parts of very different theoretical enterprises: ceteris paribus, matches light when struck; defeasibly, fish eggs turn into fish; in normal circumstances lying is wrong-making. Little and Lance think it is unlikely that there is any way to articulate, in a way that would be helpful (i.e. provide us with exceptionless principles), all of the cases in which these generalizations hold, all of the ways fish eggs turn into fish, for instance. That is, defeasible generalizations are not enthymematic for more detailed, exceptionless principles. (2008: 61) They believe that, in the way that we have generalizations that permit of exceptions in other domains we can have useful moral generalizations. In these other domains, the generalizations are meant to capture something important about the nature of their objects. (2008: 62) While it is not the case that every single fish egg turns into a fish, because, say, many are eaten by other creatures, it is nevertheless true that fish eggs turn into fish. The idea is that the claim is no less informative for admitting of exceptions than it would be if it were exceptionless. A generalization need not be exceptionless in order to be explanatory. "[The] point of the generalizations seems to involve isolating a connection that is, for one reason or another, particularly telling of something's nature." (2006b: 588)

Lance and Little claim that a central part of understanding defeasible generalizations is that we conceive of them as making reference to certain privileged conditions. The basic form of a

defeasible generalization looks something like: '(Defeasibly) All Fs are G'. For example, in privileged conditions, fish eggs turn into fish. When the generalization holds - when fish eggs do turn into fish - something is revealed about their nature. This is true even if the generalization admits of so many exceptions that it rarely holds. That is, defeasible generalizations are not mere statistical generalizations. Moreover, they not only tell us what happens in privileged conditions, when we find ourselves in non-privileged conditions, they also inform us about the ways in which our situation is different from the privileged one; the ways in which our situation would have to change in order to be a privileged one. (2006b: 588-89, 2007: 152-53, 2008: 62-63)

Take, for instance, pain. Lance and Little (2008) claim that pain is defeasibly bad-making. We learn something about the nature of pain when we learn that it is in fact bad-making. Lance and Little claim that pain-as-bad-making has conceptual priority over pain-as-good-making. Nothing needs explaining when pain is bad-making, but does need explaining when it is good-making. That is, it would be strange if, for instance, Alan insisted that the pain he experienced when Bob pulled his fingernails off with a pair of pliers were good. We would want some kind of explanation for why, given the incredible amount of pain he (presumably) experienced, he judged it to be good. But no such explanation is needed in a case in which Alan says that the pain was bad. That is just how pain works, according to Lance and Little. And to understand pain, in a practical sense, requires understanding what the relevant privileged conditions are, it requires understanding the ways in which conditions can be different from the privileged ones, and it requires understanding the differences those changes can make. (2008: 62)

Notice, however, that these requirements are requirements that any hard-line particularist, particularly of the anti-transcendental variety, would, or at least could, reject. Lance and Little are claiming that in order to have a practical understanding of the concept - the *concept* - of lying, I have to know that in privileged conditions lying is wrong-making. I have to know that when playing the card game Bullshit, I am not in privileged conditions. I have to know that the difference between being in privileged conditions and playing Bullshit is that while playing Bullshit, I and the people with whom I am playing do not owe each other the truth. But, then, this looks an awful lot like requiring a set of principles in order to understand and make judgments about lying. It is understanding the principle, 'When playing Bullshit, the otherwise binding proscription against lying is suspended.' It is understanding that in one set of circumstances the telling of an

intentional falsehood for the purposes of deceiving another is wrong-making and that in another set of circumstances it is not.

One might wonder, though, whether we need to understand how pain functions in situations other than the one in which we find ourselves in order to understand the moral significance that pain can have. Perhaps, more to the particularist's point, we might wonder whether we need to understand the moral significance that pain can have in order to understand the moral significance that pain *does* have in the situation at hand. This seems, at least in part, an empirical question about the nature of the human capacity to reason. And one thing we can take from Wittgenstein is that "Wovon man nicht sprechen kann, darüber muss man schweigen." (2003: 155)[5] To the extent that it is an empirical question, if we lack the relevant data, then what we say on the matter comes from an epistemically deficient position. But another thing we learn from Wittgenstein, which many others have noted but not quite seemed to take to its logical conclusion, comes from his so-called Rule-Following Considerations (RFCs). Put one way, very roughly, Wittgenstein's concern is that rules, e.g. for the application of a term, like 'red', or for the application of some mathematical function, like '+', are not guarantees. That is, as Marie McGinn puts it, "there is no fact about my past intention [to follow a rule for '+'], or about my past performance, that establishes, or constitutes, my meaning one function rather than another by '+'." (2006: 76) If this is right, there are no grounds for objecting to someone who, either told to or claiming to, follow the rule '+2' carries on: "96, 98, 100, 104, 108..." This is because, again as McGinn puts it, there is nothing about her past performance that establishes her meaning, as we might understand it, "+2 (forever)" rather than, again, as we might understand it, "+2 until 100, +4 after that". There is no rule to which we can appeal that would justify a complaint that she has gone on incorrectly. Translated for the moral case, the worry is that there are no grounds for objecting that someone is violating the rule she claims to follow when she tortures strangers on Thursdays, all the while claiming to follow the rule "No wanton infliction of pain upon strangers." What constitutes the rule "No wanton infliction of pain upon strangers" for her might be what we understand as "No wanton infliction of pain upon strangers except on Thursdays." And there is no way to get her to understand what we mean by the rule without appeal to some *other* rule. But then, the problem will just repeat itself, at which point we would have to appeal to another rule, ad infinitum.

[5] "Whereof one cannot speak, thereof one must be silent." (Wittgenstein 2003: 154)

Little (2000) is well aware of Wittgenstein's RFCs, discussing them in the context of the shapelessness of the moral. However, it seems that in order for Lance and Little to get their discussion of defeasible generalizations off the ground, they have to disregard the RFCs. After all, their claim is that in order to understand the moral contribution pain can make in a given situation, one has to have a grasp of the contribution pain makes in privileged conditions. Moreover, one must know that in the current situation, regardless of the contribution it actually makes, pain is still defeasibly bad-making. But all this knowledge seems to me to amount to is the possession and application of a series of rules governing the moral contributions pain makes in various situations. It amounts to the employment of a series of rules which serve to justify various judgments about what constitute privileged conditions and the various ways in which other situations can deviate from those privileged conditions. I think this point is particularly clear if we think about a dispute between two people about a particular instance of pain in which the verdicts they render contradict one another. To what could one appeal if one were in a disagreement about whether or not the current conditions of an instance of pain were privileged conditions or not? In such a dispute, there must be some articulation of what constitutes the privileged conditions of pain in order for the judgment that one of the disputants is correct and the other incorrect to be justified. And probably, if the privileged conditions reveal something about the *nature* of pain (as Lance and Little (2008: 62) say they do) then at most one of the disputants is correct. Whether some instance of pain is an instance in privileged conditions should not be just a matter of one's judgment. And if it is not just a matter of judgment, then there must be something in the world that grounds the fact of the matter. Whatever that fact, property, or relation is, it seems to me like the sort of thing that could figure in a principle articulating the privileged conditions of pain.

Wittgenstein's point is that there is more to the practice (in particular of language, and for our purposes, of morality) than the application of a series of rules. Language, as well as morality, requires sensitivity to context and judgment in order to make the right moves - in order to communicate, use words correctly in novel contexts, and also to determine which among the many actions available to one is the right action. But if Wittgenstein's point in the RFCs is correct, and can be extended to morality the way John McDowell, Little, Lance and Little, and others have attempted to do, then not only does it speak against a strongly-principled understanding of morality, but it seems to me that it speaks against Lance and Little's defeasible generalizations as well, at least

insofar as their view is still a view employing a type of rule. This would be a strange consequence of the RFCs, as the RFCs were used as impetus for Little's view to begin with. Nevertheless, it does seem that, insofar as Lance and Little's view is a view about a kind of rule - as I mention in the previous chapter, the kind of rule a particularist should look at with at least *some* scepticism - if the RFCs justified Little's first steps toward particularism, they should just as much justify a rejection of defeasible generalizations.

But where does that leave us? We want to find some kind of particularism, but all of the views we have looked at are problematic. Dancy has no response to the objections from Berker and Gert. McNaughton and Rawling, while championing the ineliminable role of judgment also have principles that are just as ineliminable as part of their view. And Lance and Little seem to have set upon a path that was either unjustified to begin with - because the RFCs do not actually support principle scepticism, or, they have circled back toward a kind of generalism that they ought to have rejected given their first steps toward particularism. So, perhaps the problem is with particularism as it has so far been conceived. In the next chapter, I will examine the prospects of some adjustments that particularists might consider making to their framework in order to get particularism to a place that cannot be co-opted by generalism in the way generalism co-opted holism. In the rest of this chapter, though, I think it would be useful to examine a view that looks quite a bit like both McNaughton and Rawling's view and Little and Lance's view, but which might (or might not be) different in important respects. One (perhaps merely superficial) difference, though, is that Pekka Väyrynen, to whom the view belongs, bills it explicitly as a moral generalist view.

4.3 Väyrynen's Theory of Hedged Principles

Pekka Väyrynen has recently developed a theory of hedged moral principles as part of a response to moral particularist claims that morality is uncodifiable.[6] Morality is uncodifiable in the sense meant by particularists and generalists if there are no useful, finite, and true moral principles. Parties to the debate generally take moral principles to articulate the relation between the non-normative properties of an action and our reasons for performing that action. In this literature, reasons are, to my knowledge without exception conceived of as contributory reasons - considerations that count in favor of performing the actions for which they are reasons. Particularists,

[6]This discussion draws mostly on Väyrynen's view as laid out in Väyrynen (2006, 2008, 2009).

again, take the holism of practical reasons to play a significant role in morality's uncodifiability. Väyrynen claims that given a few natural assumptions about the nature of normativity, he can provide explanatorily and epistemologically useful moral principles. Additionally, Väyrynen takes as one of his assumptions the holism of moral reasons. The success of Väyrynen's view, then, could mean serious trouble for moral particularism. In this section, I argue that the 'natural assumptions' Väyrynen makes about the normative landscape and the way they figure into his theory of hedged principles are too problematic to support an argument against moral particularism.

First, I will discuss Väyrynen's view, what he calls the Basis Thesis (BT) and how the Basis Thesis leads to his theory of Hedged Principles (HP). In the next section, I will describe why I think we should reject BT.

4.3.1 The Basis Thesis and Hedged Principles

In Väyrynen (2006), we are presented with The Basis Thesis (BT).

> For any consideration C that is a moral reason to Φ, the normative fact that C is a moral reason to Φ, requires a basis that explains why C is a moral reason to Φ. (719)

Suppose, for instance, that Anne, a certified EMT, sees another person, Ben, get struck in the head by a large piece of hail. Call the compound fact that Anne is trained to treat injuries and Ben is injured 'F'. Owing to F, arguably, Anne has a moral reason to help Ben should he need it. According to BT, a basis is required to explain why F is a moral reason for Anne to help Ben. So, the framework, according to Väyrynen looks like this: we have some consideration F, we have the fact that F is a moral reason (when it is), and we have an explanation (the basis) of why F is a moral reason. The basis is a normatively significant feature, e.g. a property or a relation, in virtue of which F is a moral reason (when it is). (Ibid.)

Väyrynen claims that, unless the explanation of why F is a moral reason tracks the feature (property or relation), we cannot explain the normative fact that F is a moral reason for Anne to help Ben. BT is plausible to the extent that it is the articulation of an explanatory demand that we already accept in various other domains. For instance, when a consideration is a prudential reason for someone to perform some action, we tend to think it is because of a relation between that consideration and the agent's well-being. The same is true in the epistemic domain. When some consideration is a reason to believe some fact, we tend to think that the consideration and

fact stand in some relation to each other that makes the consideration a reason to believe the fact.[7] (2006: 719-20) Because BT is a metanormative thesis, Väyrynen remains neutral regarding just what the basis (or relation) is - he remains neutral regarding just what property is that fills the normative basis role.[8] (2009: 100, 118) It is important that he remain neutral here so that he does not beg any questions about the truth or form of moral principles. Väyrynen's strategy is first to argue for a general view about the way reasons work which is neutral between particularism and generalism. He then argues that BT allows us to make the best sense of the apparent holism of reasons when conjoined with a theory of hedged moral principles. Because his discussion is a metaphysical one about both bases of reasons and the principles in which those reasons figure, and not about how we get to know or work with them, his position is still neutral between particularism and generalism.[9] As Väyrynen notes, it may well be that principles are part of the furniture of moral reality without it being the case that we rely on them in a way that would be objectionable to some particularists. (2009: 111) After all, Dancy claims that holism leads to particularism. If BT is consistent with holism, then it ought also to be consistent with particularism. Indeed, more than being consistent, BT helps us to make sense of and explain the apparent holism of reasons, so it should help us to explain particularism if holism *does* indeed support particularism.

Consider a simple, silly case to illustrate holism. For instance, consider a case in which we must assess whether there is a reason that favors lying. Sometimes we have reason to lie and sometimes we do not. Suppose I spent several hours preparing and cooking a meal for my wife and it turns out to be disgusting. Now, among the various responses she might offer when I ask her how she liked the meal, she finds herself trying to decide between two: either she lies to me and tells me that the food is not bad or she tells me that it is disgusting. Honesty is important, so the lie might seem, for that reason, pro tanto wrong. But not hurting one's spouse's feelings is also important, so perhaps, to that extent, the lie would not be pro tanto wrong. Indeed, while it seems like it

[7]It is worth noting that, given the ways in which the reasons in various domains differently relate to things like truth and rationality (in those domains), drawing analogies across normative domains like this may not be as useful as it is alleged to be. Again, see Gert (2008) for a nice example of an argument against the presumption that the features in the epistemic domain to which theorists appeal in their arguments for theoretical reasons' holism appear in the practical domain.

[8]I say that Väyrynen remains neutral regarding just what property is that fills the basis role and not 'just what the property or properties are that fill the basis role', because as Väyrynen articulates his view, and as I go on to discuss, he does seem to think that for each moral action-type, there is a single property fulfilling the basis role. By 'moral action-type' I mean something like 'all actions of the type: (lie, cruel, kind, just, etc)'.

[9]Specifically, it is neutral about the truth of Anti-Transcendental Particularism - Dancy's view according to which the possibility of moral thought and judgment does not depend on a suitable supply of moral principles.

would be permissible for her not to lie to me, it would also be permissible for her to lie to me. But clearly in other cases it would be wrong to lie. For instance, suppose I have a piece of spinach in my teeth and I ask her if I am clear of spinach before an important meeting. In this case, I take it that lying would be wrong. According to Väyrynen, what explains why the fact that she would be lying to me makes my wife's action pro tanto wrong is the fact that in the one case, and not the other, the normative basis of the wrongness of lying is instantiated.

Ultimately, the goal in Väyrynen (2006) and (2009) is to show how we can get explanatorily and epistemologically useful moral principles just from examining a plausible way the normative domain might work. In order to avoid the morass surrounding the atomism versus holism debate, Väyrynen just grants holism and tries to show how, even if true, there is still room for principles. Traditionally, debates between particularists and generalists over the 'right kind of principle' have been debates about finite, exceptionless, true principles. But many now believe that such principles are not really the proper subjects of the debate. And so, this sort of principle is not what Väyrynen is after. Indeed, as he notes, very few ethicists are looking for the kind of exceptionless principles that used to get bandied about in the literature. (2009: 92-93) No one, perhaps apart from Kant and his strictest adherents, thinks that principles of the form, "Lying is wrong" are exceptionless. So what Väyrynen is after is a kind of principle that is true, is manageably usable by creatures like us, and helps us to explain why and how the things that are right and wrong are right and wrong.

First, for clarity's sake, here I will reproduce Väyrynen's "basis argument" for moral generalism which begins with BT as an assumption.

> (G1) Any moral reason requires a normative basis that explains it.
>
> (G2) The normative basis of any moral reason requires the existence of a (set of) true moral principle(s).
>
> (C1) Therefore, moral reasons depend for their existence on the existence of a comprehensive set of moral principles.
>
> (G3) The possibility of sound moral judgment depends on the existence of moral reasons.
>
> (C2) Therefore, the possibility of sound moral judgment depends on the existence of a comprehensive set of moral principles. (2006: 722)

Väyrynen claims that, because BT is compatible with particularism, G1 does not imply anything about a connection between moral principles and moral reasons. Likewise, we ought to accept G3

because if it were false, our judgments about moral reasons would be false, so we have to grant G3. This, so far, is only slightly objectionable. Below, I will question whether BT is, in fact, compatible with particularism. Particularists ought to reject G2. The move from G1 and G2 to C1, if sound, would be very damaging to the particularist project. It seems to me, if Väyrynen gets C1, then he has already done enough to do away with particularism. In what follows, then, I will attempt to stop Väyrynen getting to C1.

So we started with BT, the claim that something - some factor or relation - is necessary to explain what makes a consideration the reason it is when it is such a reason. BT also allows us to explain why, if that same consideration is not a reason in other situations, it is not the reason it normally is. It is useful when discussing schematic views to employ substantive examples to illustrate the ways in which the schema works. To that end, and following Väyrynen, I will use a principle about lying as my candidate principle. Väyrynen suggests two candidate theories articulating the normative basis of the wrongness of lying - though he does not commit himself to either. Because neither I, nor Väyrynen, are committed to either of these candidate theories, and for brevity's sake, I will only discuss one of them.

One potential reason why telling lies is wrong is because it undermines a useful social practice. Given Väyrynen's view, then, in the above case involving the disgusting meal I made for my wife, what makes it the case that lying to me is not pro tanto wrong is the fact that that lie does not instantiate the normative basis of the wrongness of lying. It is not the case that the lie would undermine a useful social practice because, arguably, a social practice that required we tell the truth even when doing so would hurt each other's feelings without any compensating benefit would be a harmful one. And similarly, the reason why lying about whether or not I have spinach in my teeth would be pro tanto wrong is because, arguably, this lie does undermine a useful social practice. Being able to ask those around you if you have something in your teeth is a lot more convenient than having to find a mirror. It saves us time and energy. As I said, this a rather silly case. Nevertheless it helps to illustrate the point Väyrynen wants to make - that the notion of a normative basis (whatever it is) helps us to explain the way reasons seem to function.

Reasons are one thing, principles another. A brief survey of the literature shows how parties to the debate frequently move from discussing reasons to discussing principles and back again, sometimes without acknowledging the shift. Väyrynen, though, tries to make it clear that BT fits

into a theory of how moral principles work, as hedged principles. "Substantive principles concerning moral reasons can be captured by principles which are hedged by reference to the normative basis of those reasons." (2009: 104-05) So we are then offered a principle articulating the wrongness of lying:

> (Lie) Something's being a lie is always a reason not to do it, provided that it instantiates the designated normative basis for this fact's status as [a] moral reason not to lie. (2009: 105)

Again, Väyrynen does not articulate what the normative basis is here because it is not necessary to identify that property in order to show that his structural account works. Of course, this means that until we *do* pick out a property to play the normative basis role, we cannot identify which considerations provide reasons against lying, nor which lies are permissibly exceptional. We can, however, generalize (Lie) to give us the form hedged principles will take:

> (HP) Any x that is G is M (e.g. x's being a lie is always a moral reason not to do x), provided that x instantiates the designated normative basis of G's contribution to M. (2009: 106)

Väyrynen is quick to note that simply the articulation of (HP) does not imply that there are moral principles, nor does it imply that any moral principles that there may be permit exceptions. For all Väyrynen has said, (HP) might be the best way to articulate what moral principles must be like, but there might nevertheless not be any moral principles. What it does do, though, is allow for the possibility of exception-permitting principles. And, again, how this works out precisely is determined by the property selected as that which fills the designated normative basis role. For Kantians, that property will have something to do with rationality and autonomy. For Utilitarians, that property will have something to do with the distribution of utility.

Whether we can find the properties that play the role of designated normative basis for the various moral principles that moral theorists want to employ in their theories is still an open question. Väyrynen thinks that justification in ethics would be arbitrary if we did not have some explanation as to why the considerations that are reasons in a situation are the reasons they are and why, if they are not those same reasons in other situations, they are not. I will raise questions about the concept of a designated normative basis below. For now, though, I shall assume that there are no problems with finding the designated normative basis of a given reason.

If Väyrynen is right up to this point, then we should clearly be able to see how his hedged principles are explanatorily useful. A given principle explains why some act is wrong when it is because the instantiation of the property playing the basis role articulated in the principle is flagged by the principle as bearing a relation to the action's deontic status. The principle can also explain why some act is permissibly exceptional - that is, though normally wrong, not wrong here (nor a counterexample to the principle) - because the principle articulates the relation between some property whose instantiation is necessary for the action in question to have a particular deontic status, and the absence of that property indicates that the action does not have the deontic status it otherwise might. This explanatory role is an important one that principles are meant to, and ought to, play. We might wonder, however, whether hedged principles as they are described by Väyrynen can play the kind of epistemological role that moral agents need them to play if the principles are to be useful. That is, Väyrynen's hedged principles seem not to do a much better job of providing moral guidance than traditional principles. I go on to discuss this further below.

4.3.2 Usable Hedged Principles

Now, part of the impetus for the entire debate between particularists and generalists has to do with the fact that, if moral reasons vary with context, then on the classical, exceptionless understanding of principles, the principles subsuming those reasons will admit of exceptions. The objection is that if a principle admits of an exception, it is false. Väyrynen is trying to show that there can be a type of principle that admits of exceptions but is not false. We see these sorts of principles (or generalizations) in other domains of inquiry. For instance, the claim 'Fish eggs turn into fish', a claim we generally admit is true despite the fact that the majority of cases are exceptional, relies on the instantiation of the property to which fish eggs owe their turning into fish. When this property is not instantiated, we have, perhaps, caviar. Väyrynen tells us that his account of hedged principles suggests an account of what it is to judge a case as permissibly exceptional. (2009: 95, 96ff.) According to this view, whether or not we judge a particular instance of, say, lying, as permissibly exceptional or not depends on our conception of the normative basis of the wrongness of lying. Väyrynen, though, admits that our knowledge of the kinds of cases in which the normative basis might be instantiated is incomplete. This is not a problem for his view,

as we can nevertheless be guided by the way we think about lying and the normative basis of its wrongness.[10] (2009: 119)

According to Väyrynen, as long as I think that there is some normative basis of the wrongness of lying, then a hedged form of the principle against lying is available. He claims that even if my knowledge of the normative basis is incomplete or incorrect, this is not the kind of incompleteness or incorrectness that causes a problem for the principle. Even if my knowledge of the property is incomplete, my conception of it may nevertheless be complete enough for me to be guided by it, at least, in a range of cases. (Ibid.) For instance, though I may not know all of the varieties of permissibly exceptional lies, I might believe a certain class of 'white lies' to be inconsequential enough to be guided, and guided correctly, in a host of cases in which nothing important hinges on the truth being told. As Väyrynen says, what guides me is my conception of the principle.

> Typically our acceptance of a principle like 'Curtailing freedom of expression is [pro tanto] wrong' is not brute. I would be a defective moral agent if I thought, for instance, that it is wrong for the government to censor the press or ban protests at speeches by its officials, but didn't think that there was any basis for judging such government actions to be bad. (Ibid.)

So, at least typically, my being a competent moral agent depends upon my judging there to be a normative basis of the reasons for and against the actions I judge or perform.

But as competent moral agency depends upon one's being a competent judge of normative bases, might one's competency judging normative bases depend upon some other judgmental capacity? Perhaps, there needs to be a basis for every normative basis, and one is competent in judging the latter only if one is competent in judging the former. And quickly, a regress ensues. Väyrynen (2009, 2006) briefly address the concern about a regress of normative bases and the judgments required to recognize them. Väyrynen maintains that the normative basis of a feature's being a reason need not be distinct, in some sense, from the feature itself. The feature need not be conceptually distinct from the basis if the feature is, for example, a thick moral concept, like justice.[11] Part of the concept of justice is that to the extent that some action is just, one has a reason to do it. Nor need the feature be metaphysically distinct. For instance, Väyrynen asks us to suppose that the

[10] Perhaps in the way one might be guided by the conception of a speed limit without knowing what the actual speed limit is at a given time.

[11] Precisely how to work this out will depend on one's conception of a 'thick concept'. I will leave such questions to one side.

fact that something promotes well-being is a moral reason to do it. Then, if the property of being a moral reason is reducible to the property of promoting well-being, then the feature of the action and the basis are not distinct.

So, it looks as though if we grant BT, then employing hedged principles is one way of making sense of the the alleged generalist nature of normativity generally, and morality in particular. From BT and hedged principles we get an account of how reasons seem to work holistically, we get explanations of why certain actions are sometimes right and sometimes wrong, and we get a suggested account of how hedged principles figure in our moral thinking. In the rest of this chapter, I will argue that, though some variety of particularist can accept Väyrynen's view, we ought to reject it.

4.3.3 Against BT and HP

Up to this point, for the sake of the argument, I have accepted BT. However, I think there are a handful of problems with it that merit further attention. My hope is that the force of the problems I raise will be stronger than the intuition that motivated BT, which is a particular instance of what Väyrynen (2013) calls 'the dependence intuition'. If the problems have more force than the intuition, then perhaps, whether BT is consistent with particularism or not, we ought to reject it. If not, however, then I will turn my eye toward Väyrynen's conception of hedged principles to examine whether or not they are the kinds of principles that can cause problems for particularism. If they are, then the particularist needs a reply to G2, the claim that "the normative basis of any moral reason requires the existence of a (set of) true moral principle(s)." By the end of this section I hope to have done enough to show that at least one of G1 or G2 of the Basis Argument must be rejected. Recall, G1 is just an abbreviated version of BT - the claim that moral reasons require normative bases. (2006: 722)

One concern regarding the basis thesis is whether for a given action-type[12] the designated normative basis must be the same for every reason not to perform that action. Väyrynen writes as though the designated normative basis is the same when discussing what makes an exception permissible. (2009: 96ff.) For instance, he proposes two toy theories to explain when a lie is permissibly exceptional. He suggests that one theory might claim that being the telling of a lie is a

[12]Again, by action-type, I mean all actions of the type: (lie, cruel, kind, just, etc.). For instance, all lies belong to the moral action-type: Lie. All just acts belong to the moral action-type: Just. Etc.

reason against performing an action "when, and because, lying contributes to undermining such beneficial social practices as trusting other people's word." (2009: 97) This account of the designated normative basis of the reason against lying would fit well as part of a variety of normative ethical theories. For instance, this sounds like something a rule utilitarian might accept. Remember, though, up to this point - through his discussion of what the designated normative basis *is* - the view Väyrynen is presenting is not supposed to be inconsistent with particularism. It seems, though, that if its undermining such beneficial social practices as trusting other people's word is the designated normative basis of reasons not to lie, then we have already got an objectionable (from the particularist perspective) moral principle - viz. 'one ought not to perform actions that undermine beneficial social practices'. Perhaps, though, this is not problematic. Väyrynen is offering a structural account. It is up to us to fill in the property that plays the normative basis role. Väyrynen's view is not utilitarian or deontological, rather it is only a metanormative framework. Nevertheless, it looks as though whatever property we inject into the framework to play the normative basis role we end up with something already looking like a moral principle. But that, in a way, just is Väyrynen's point. When we talk about moral reasons, we cite the reasons' normative bases. In picking out the normative basis of a reason, we have got a particularized version of a hedged principle. For instance, the (moral) reason why it is wrong to lie (here), is because telling the lie would undermine a beneficial social practice. And, insofar as we find ourselves already talking about hedged principles, just in virtue of citing a reason, certain particularists ought to wonder whether they really can take on board the basis thesis.

Now, if someone is asked on multiple occasions why certain instances of lying are wrong - that is, if someone is asked what the normative bases of certain lies' wrongness are - she might provide different answers on those different occasions. These different answers might reflect the fact that she thinks differently of the reasons why the lies are wrong. This might be because, given her theory, the normative basis really does change. Alternatively, it might just be that she cites the most salient, but derivative reason on various occasions. In that case, if we were to work out her theory, we might have, at bottom, one consideration that manifests itself in different ways in different situations. However, if it is possible for what Väyrynen is calling "the" designated normative basis to be just one among *many* such normative bases for a single type of reason (so, X, Y, Z might be the bases for different instances of lying), then, it seems to me, what has happened is that Väyrynen

has not provided a viable solution to the problem particularists have with generalism. Instead, he has just changed the locus of the problem.

Part of the problem early particularists raised for generalism was that it is difficult to see how we can be confident in extending our judgments from one moral case to future moral cases. So, for instance, there is nothing in the combination of properties A, B, C, and D and our judgment, from those properties, that we ought to Φ that entails that every time those properties are together we ought to Φ. While in one situation we might have A, B, C, and D, in another situation we might also have E, which changes the way A, B, C, and D combine and so warrants a different judgment about Φ-ing. This just is holism - which Väyrynen accepts. Väyrynen is saying that we can have these hedged principles and a conception of the normative bases of various reasons, and together they can help us to make moral judgments across contexts. This is because we can know that if the normative basis is instantiated, then the reason functions in a particular way, and if the normative basis is not instantiated, then the reason does not function the same way. But note that if it is possible for the normative basis in one situation to be different for the same type of reason - a reason against lying, for instance - then in order to judge whether we are looking at a permissible exception, we have to be in a position to know, or to have a conception of, the various different normative bases for the wrongness of lying. And that presents the same sort of epistemic problem that we have seen before. We might not be in a position to know all the ways in which various reason-giving considerations can come together to yield verdicts about what we have most reason to do. So, despite BT and HP, maybe the ineliminable role for judgment that Dancy, McNaughton and Rawling, and others indicate is where particularism maintains its foothold in the debate is still there.

Suppose, further, that we do have to be able to identify the various normative bases for a particular reason - again, say, a reason for or against lying. What this means, given the way BT features in hedged principles, is that we would have an explosion of moral principles - perhaps one for every normative basis. We would have to have, in our possession, a principle for every way that a lie can be wrong and for every way the lie can be permissibly exceptional. The same would be true for every other moral action - killing, stealing, sharing, being kind, etc. Indeed, it seems to me that the mere potential of an explosion of moral principles of which we would have to be in possession means that the human capacity to be aware of normative bases has to be open-ended. That is,

because we cannot rule out the possibility of novel permissible and impermissible exceptions, then to the extent we want to be conscientious moral agents, we ought to be skeptical that the moral principles of which we are in possession tell the whole, or perhaps even a sufficient amount, of the story.

Väyrynen (2006) considers the objection that multiple normative bases for a single type of moral reason would be problematic. He offers two potential avenues of response. My earlier example in which someone is asked on multiple occasions whether an instance of lying is wrong mirrors these response. In the first instance, he says that the fact that distinct normative bases would yield distinct principles is no problem at all. "But the resulting proliferation of principles would be unobjectionable, since the [n] principles would reflect distinct moral concerns...and have distinct application conditions." (2006: 734) The second reply he offers is that, in fact, at most one of the principles on offer is true. The bases of the other candidate principles are relevant in some other way to the wrongness of the action, but not as wrong-makers themselves. So, for instance, if the true hedged principle covering lying cites its detriment to a beneficial social practice as the normative basis, then a candidate principle which cites the fact that the lie undermines someone else's autonomy is merely citing an intensifying consideration. Such a lie would strengthen the reason not to lie, but would not itself provide a reason against lying if the autonomy-underminingness of it were the lone consideration.[13] (Ibid.)

Väyrynen's second response looks like it amounts to a denial of what we might call 'normative basis pluralism'. This would be the denial of the possibility of multiple bases being the bases for a single type of moral reason. In a sense, the claim is like some remarks McNaughton and Rawling make about can-openers. At a low, physical level, there is nothing that clearly unifies this collection of properties or this collection of considerations. What unifies them is at a higher, functional level - that they are all can-openers. Similarly in the case of moral reasons, what unifies this collection of very disparate normative bases is the fact that they are all the normative bases of, say, reasons against lying. Or, at least, that *could* be a reply to the objection if Väyrynen's second response accepted the possibility of disparate normative bases.[14] Instead, though, he denies this and claims that at most one of the normative bases is the correct one. Any other consideration that

[13]Väyrynen's example in this context involves killing. For continuity's sake, I have here changed it to lying.

[14]In fact, in personal correspondence, Väyrynen claimed that he is happy to allow basis pluralism. However, it is unclear how he can do this given the claim he makes in Väyrynen (2006) that at most one of the bases is *the* designated normative basis, while the other alleged bases are relevant only in a non-reason-giving fashion.

is morally relevant is not relevant as a wrong-making feature of lying. (Ibid.) As a particularist, it seems to me obvious that there can be a variety of explanations for the wrongness of a series of lies. This by itself is not an objection, of course. As a particularist, I am starting from a very different place than Väyrynen is in his theorizing. However, there seems to be some tension between Väyrynen's position, i.e. that of adopting holism and attempting to start from a place unobjectionable to particularists, and the apparent claim that for every moral action-type there is only one normative basis shared by all actions of that type. The particularist commitment to holism is to a thoroughgoing holism. It is unclear why that holism should not extend to a holism of normative bases. If that is the case, then Väyrynen's holism is not that of particularists and it is unclear whether his view does start from a neutral place.

Perhaps humor can provide a relevant and insightful analogy. Whether a particular joke is appropriate is highly context-dependent. On the model provided by BT and HP, when a joke is appropriate we might say it is because the normative basis of its being inappropriate is not instantiated and when a joke is inappropriate the normative basis is instantiated. But of course, the reasons why a joke might be inappropriate in one context can be dramatically different from why it is inappropriate to tell in another context. Some jokes are inappropriate to tell - that is, we have reasons not to tell some jokes - because we do not know the target of the joke - she is a stranger, for instance. Other times, it might be inappropriate to tell that very same joke because in doing so, we would be bringing to light something we were told in confidence. Still at other times, it might be inappropriate to tell the joke because doing so would just be mean. These are different reasons why the inappropriateness of a joke gives us a reason not to tell it. And it seems to me the same basic line can be taken with a more obviously moral case involving, for instance, lying.

Consider now the first response Väyrynen offers to the worry that there might be multiple normative bases for a single moral property. He claims, roughly, that if there are different normative bases for the same moral properties, then this will just yield distinct principles covering, say, lying. But, Väyrynen claims, since the principles covering lying's myriad normative bases would be distinct principles with unique application conditions, there is nothing objectionable about them. (2006: 734) But this reply seems just to turn the set of principles covering lying into something like a multiply disjunctive exceptionless principle covering lying. One objection to traditional principles is that we would have to build into the principle all of its exceptions. If we did this, and the

particularist's conception of the complexity of the moral landscape and the myriad ways in which features can modify reasons is accurate, the principle would be so long that we could never hold it all in our minds at one time. Väyrynen's acceptance of myriad distinct principles covering lying might render the set of principles covering lying similarly large and complicated. This, by itself, may not be a problem. But when combined with the fact that, in the midst of a decision about whether or not to tell a particular lie, we (regular humans) ought to be able to consult the relevant principles, it might be a problem. It is probably plausible that sometimes we are not aware of all of the relevant moral considerations in a situation. Sometimes we have to examine the situation and weigh considerations against one another. The task of deciding whether or not to lie seems to become increasingly unmanageable the more principles we have to consult to determine whether such and such a consideration is the normative basis here or whether it is trumped by another consideration. If this is a consequence of Väyrynen's view, then it is objectionable and not just on particularist grounds.

If the preceding objections to BT and HP are right, then we have a dilemma. Either for each type of moral reason (for all reasons against lying, say) there is one normative basis or there are different normative bases for the same type of moral reason. If the former, then it seems we have ruled out a kind of pluralism that might reasonably be viewed as an important part of a variety of particularisms. If the latter, then it seems that the principles covering the reasons we have might be more numerous than we can handle.

All this said, though, the intuition motivating BT is strong. Perhaps there is some way around these worries, or perhaps they are not in fact worries at all. Perhaps owing to my own particularist intuitions, I cannot view the matter rightly knowing the problems that Väyrynen's employment of BT causes particularism. Perhaps the only way to rescue particularism from Väyrynen's theory of hedged principles is to let BT pass, and attack his conception of hedged principles directly. In the rest of this section, I attempt to take down hedged principles while granting BT.

Recall, a hedged principle is hedged because of the built-in reference to the normative basis, which may or may not be instantiated in any particular instance. So a principle articulating (or capturing) the wrongness of lying would say that lying is wrong when and because the property playing the basis role for the wrongness of lying is instantiated in the situation. The same thing, mutatis mutandis, is the case for other things that are either right or wrong, like killing, stealing,

torturing, generosity, kindness, justice, etc. So, for instance, that Φ-ing is cruel is a reason not to Φ when and because the property which plays the basis role for the wrongness of cruelty is instantiated by Φ-ing.

But how does knowing the alleged truth of a hedged principle help us to navigate the moral landscape? Do hedged principles do better than the classic (straw-man?) conception of an exceptionless principle? One thing we seem clearly to want in a moral principle is some sort of guide - it ought to help us to make our judgments and decide which actions to take. It is an interesting, though it seems empirical, question to what extent our conception of a moral principle *actually* guides our actions. How much of the situation do we have to identify as morally relevant in order to know whether a particular principle applies? Do we have to know the principle in addition to knowing the relevant features of the situation? Or is it instead a matter of having the principle in the background, in our minds, before we are in a situation which enables us to see the morally relevant features as morally relevant? At which level does the principle operate? That is, are principles like the rules of a game, which we consult in a rulebook? Or, are they something more fundamental? Do moral principles work in the background of moral perception the way edge-detection works in the background of visual perception? This is more than just a series of rhetorical questions. These are substantive questions to which any decent generalist account ought to provide, or at least suggest, answers.

It does not seem to me that Väyrynen's hedged principles answer these questions much better than traditional, exceptionless conceptions of moral principles. Suppose the designated normative basis of the wrongness of lying is that it has the property, S, it undermines a beneficial social practice. What seems to be important for Väyrynen's view is not that this is a lie. Rather, what is important is that it is an action that has this normative property - which Väyrynen is calling the designated normative basis for the wrongness of lying. It does not matter that the action is the telling of a lie. What matters is that it is an action that has S. Presumably, S can be instantiated by other actions that are not lies. And, presumably, someone according to whose theory the wrongness of lying is due to lies having S would suggest that other actions that have S are also, to that extent, wrong. But then, are the properties of actions such as 'being a lie', 'being unjust', 'being cruel', just idle cogs? Given the fact that particularists are meant to be able to take BT on board in conjunction with their commitment to the hyper-complexity of the moral landscape, particularists

might be able to talk about BT in conjunction with their talk of reasons but pull up short of talking about principles, precisely because all of the work can be done by reasons-in-situ. Principles would then be superfluous.

Väyrynen (2009) responds to the worry that principles are idle cogs once we have a grasp of normative bases. (115ff.) He says,

> The designated normative basis is just a property which particular lies instantiate or not. It exhibits systematic patterns of counterfactual dependence between whether something is a lie and whether there is a moral reason not to do it only when embedded in a generalization like (Lie). What (Lie) asserts is precisely a complex but systematic relationship of dependence between these two factors and the designated normative basis. It asserts a connection between something's being a lie and there being a moral reason not to do it which is stable under any hypothetical changes under which it still instantiates the designated normative basis, but which might not hold outside this range of conditions. So the designated normative basis explains moral reasons in a systematic way only given a principle like (Lie). (2009: 115)

It is not at all clear to me that this response adequately addresses the worry. The complaint I am currently making against Väyrynen's view is that once we have a handle on what the relevant normative basis in a situation is, we do not need a principle like (Lie) which cites the designated normative basis as being the basis of some other feature's normative status. Väyrynen claims that (Lie) is important because (Lie) draws the connection between the designated normative basis of a particular lie and the moral reason not to tell that particular lie. But it cannot be the case that we cannot understand the basis relation without appealing to some generalization like (Lie), because then BT, having a necessary epistemological role for principles as a *prerequisite*, would beg the question against particularists. And recall, BT is supposed to be neutral between particularism and generalism. So, particularists ought to be able to understand the connection between a particular lie and the normative basis of the wrongness of lying without appeal to a generalization like (Lie). But then, again, (Lie) is superfluous - the alleged explanation of the systematic patterns of counterfactual dependence notwithstanding. Indeed, for a particularist, the alleged systematic patterns are either merely alleged or themselves superfluous. The particularist position is just that these systematic patterns are, in fact, not there, or that they offer no necessary assistance to us in our efforts to figure out what we have reasons to do.

If the preceding arguments are roughly correct, then it looks like the move from BT to HP is unnecessary. However, given the concerns raised about BT, Väyrynen may not even get as far as being in a position to attempt that move. One way to respond might be to claim that, for the sake of argument, we have to suppose that the concept of a reason with which particularists are working is coherent - despite the fact that it is not. That is, the concept of a reason may, at least implicitly, appeal to general rules. As McNaughton has put it, moral principles are just reasons that have had their generality made explicit. (1988: 191) The concept of a reason, that is, may just imply an appeal to some generalization.[15] The response, then, is that to the extent that particularists deny this implicit appeal to generality, they employ an incoherent concept of a reason. It is not the fault of BT or HP that particularists are working with incoherent notions. So, the fact that, despite his best efforts to make it neutral, BT still favors generalism over particularism is just a consequence of how normativity works, and not because of any questions begged by BT.

At this point in the dialectic, though, we seem to come to an impasse. We come to the root of the debate between particularists and generalists. Though both particularists and many generalists take reasons to be the fundamental normative units, they disagree, inter alia, about how reasons work, what they are, and how they add up together. Until these metanormative questions are solved and while particularists and generalists maintain that reasons are the basic normative units, they will continue to talk past one another.

In this chapter I have discussed two moderate generalist views as potential alternatives for particularists. The thought was that, if particularism has internal problems as has so far been the case, then maybe continuing the trend of moderating the view (i.e. conceding ground to generalism to make particularism less jarring) would lead us to an acceptable middle-ground between strong particularism and strong generalism. While there are very few unproblematic and substantive views in philosophy, Lance and Little's defeasible generalizations and Väyrynen's hedged principles require too much of particularists. Particularists who already accept the strength of Wittgenstein's RFCs cannot accept either defeasible generalizations or hedged principles. In the next chapter, I suggest that what particularists ought to do is to go back to the other side. Particularists must change the state of play in order to defend the intuition that moral principles fail to tell the whole story. To that end, I will be suggesting that particularists reconsider their conception of a practical

[15] Again, see Gert (2007) and Berker (2007) for discussions of how, because particularists lack this implicit appeal to generalizations, their conception of a reason is allegedly incoherent.

reason. What once was the primary weapon in the particularist arsenal - holism about reasons - must be abandoned, or at least heavily modified.

CHAPTER 5

DOING WITHOUT THE CONTRIBUTORY

5.1 Introduction

Let's review for a moment where we are at this point. In Chapter One, I laid most of the groundwork for the presentation of the theories that followed in Chapters Two through Four. In Chapter Two, I discussed the particularist view presented by Jonathan Dancy - his Anti-Transcendental Particularism. This view maintains that the possibility of moral thought and judgment does not depend on a suitable supply of moral principles. (Dancy 2004: 7) I also presented what I take to be some of the most serious objections to Dancy's view - those proposed by Sean McKeever and Michael Ridge (2006), Joshua Gert (2007), and Selim Berker (2007). McKeever and Ridge claim that Dancy's view is problematic in at least two ways. First, his conception of a default reason is incoherent. Second, his claim that holism lends support to particularism is false. In the first instance, default reasons are problematic because Dancy claims that a default reason is one that comes already 'switched on' - either favoring or disfavoring some action. When this happens, there is nothing about the reason that needs explaining. For instance, when the fact that Φ-ing is a lie counts against Φ-ing, because the default status of lying is such as it is, we need not explain why Φ-ing is pro tanto wrong. But in cases in which the fact that Φ-ing is a lie does not count against it, or in a case in which the fact that it is a lie counts in favor of it, there is something else present (or missing) from the situation that makes the lie function differently from the way it does by default. The problem with this account is that Dancy relies upon his distinction between enablers and disablers for his explanation of default reasons. McKeever and Ridge aptly note that it is difficult to distinguish between the presence of an enabler and the absence of a disabler. Dancy's default reasons rely on this distinction, but because the distinction itself is untenable, so too is his concept of a default reason.

In the case of Dancy's claim that holism supports particularism, McKeever and Ridge note that generalists can build principled theories from frameworks employing a holistic conception of reasons just as easily as particularists can. If they are right about this, then it seems as though

particularists' main expressed motivation for their view fails actually to lend their view any support. And in Chapter Four, we saw two examples of non-particularist views which do take holism on board. The potential legitimacy of such views might be problematic for particularists who take holism to provide unique support for particularism.

Gert (2007) and Berker (2007) independently developed objections that are similar to one another. They both focus on the extreme context-dependence of Dancy's holism. Their claim is that Dancy's concept of a reason for action is incoherent due to his inability to explain what it means for a reason to have any kind of (stable) strength value. Their arguments point to the fact that if reasons are as context-dependent and sensitive as Dancy says they are, then we have no way of talking about the strength of an individual reason. This is because the only experiences we have with reasons are *in situ* - where their strength is influenced by the other features of the context. And as Dancy has argued elsewhere, trying to pull the reason out of context will not help us to figure out what it does in a context. (2006: 40ff.)

I went on in Chapter Three to discuss the particularist view put forward by David McNaughton and Piers Rawling. According to their view, which they call Thick Intuitionism, there are some principles - weak principles - that frequently involve thick moral concepts. Because these weak principles involve thick concepts, they are already normative. That is, McNaughton and Rawling take a (perhaps *the*) central aspect of particularism to be a rejection of the possibility of so-called 'bridge principles'. Bridge principles, in this context, are principles connecting the wholly non-normative with the normative. So, while we can rely on some principles, the moral landscape is not quite as principled as some generalists claim it is. For instance, moral situations are too complex and complicated for principles such as the classic utilitarian principle that the right action is the one that produces the greatest balance of pleasure over pain to be of any use. I suggested that insofar as McNaughton and Rawling accept and allow a role for so-called weak principles in their theory, they come far too close to being out-and-out generalists. Indeed, as I say in that chapter, McNaughton admits that their view is, at a certain level, fully generalist. Whether the view is without internal problems or not (and it is not clear to me that it is free of them), to the extent one is looking for a fully particularist view, their view is not a viable option. And as I discussed in that chapter, one reason for looking for a 'fully particularist' view is that it seems like one way, and to

my mind the best way, to avoid the many problems that generalist views and extant particularist views have is to move further away from the generalist end of the spectrum of views.

Because of the apparent problems with the two best-known varieties of particularism - Dancy's and McNaughton and Rawling's - I considered in Chapter Four two other views that might appeal to theorists who have sympathies with the particularist project but agree that Dancy's and McNaughton and Rawling's views are problematic. In the face of objections raised by generalists, the quintessential particularist maneuver has been to modify the particularist position, effectively appeasing generalists.[1] I considered the moral contextualism offered by Margaret Olivia Little and Mark Norris Lance. Their view, which grew out of Little's (2000) particularism, agrees with particularists that there is not much hope of finding true, useful, exceptionless principles. However, they disagree with particularists that there are no true, useful, *exception-permitting* moral principles. Instead, Lance and Little argue for a type of defeasible generalization - akin to generalizations we find in the sciences. Their view is that just as the claim 'Ravens are black' is true, despite counter-examples, and useful because of what it tells us of the nature of ravens, so too can a claim like, 'Lying is wrong' be true - despite counter-examples. Claims like 'Lying is wrong' are useful because they tell us something about the nature of lies. Lance and Little's defeasible generalizations are meant to respect the complexity of the normative landscape while at the same time admitting that there seems to be some uniformity to the landscape at some level.

The next view I considered was the theory of hedged principles developed by Pekka Väyrynen. Väyrynen argues that, given a few natural assumptions about how normativity works, one can develop an explanatorily and epistemologically useful theory of hedged principles. This theory of hedged principles can offer, according to Väyrynen, generalist responses to the concerns about morality that particularists raise. For instance, particularists frequently point to the complexity of normativity and the variability of reasons. When particularists point to the alleged truth of holism as providing support for their view, it is because they think that the variability of reasons leads to the uncodifiability of morality. They are pointing to the fact that for any rule about the way a reason works, there are exceptions that cannot be captured by a principle that is finite, useful, and true. In articulating the principle, one of those three features - finitude, usefulness, or truth - has to be sacrificed. What Väyrynen's hedged principles purport to provide is an account of how we

[1] We see this in the shift from Dancy (1993) to Dancy (2004), in which Dancy goes from talking briefly and off-handedly about default reasons to devoting a large discussion to the the topic.

can have a finite, useful, and true (true enough?) principle. Like Lance and Little's view, he draws an analogy between his hedged principles and generalizations offered in the sciences. 'Ravens are black', 'fish eggs turn into fish', and 'acids are corrosive' are all claims that admit of exceptions - there are albino ravens, some fish eggs are eaten, and some acids do not corrode the things to which they are applied. But, the argument goes, the fact that these claims admit of exceptions does not mean that they are not true. What it shows is that, in most instances, the purported exceptions are not relevant to the truth of the particular claim. For instance, when one claims that ravens are black, there is meant to be an implied bracketing of albino ravens. The fact that albino ravens exist is not relevant to the truth of the claim that ravens are black because albino ravens are aberrations.

By way of response to these two views, I claim that, in the first case, Lance and Little are subject to a similar objection raised by McKeever and Ridge against Dancy. Given that their view started as particularist, citing Wittgenstein's Rule-Following Considerations, the fact that they end up with a theory that is subject to an objection based in the Rule-Following Considerations is a bit ironic and problematic. Moreover, and more importantly, it seems that in order to employ defeasible generalizations in the moral domain, we would already have to have a solid grasp of the cases in which the alleged principle applies and those in which it does not - which seems to presuppose an understanding of the guidance that moral principles are supposed to provide us to begin with. Since this is precisely one of the points of contention, it cannot be the point from which Lance and Little develop their view.

Against Väyrynen's hedged principles, I argue that the 'natural assumption' about normativity - that for any reason we need an explanation of why it is the reason it is - causes problems. I also argued that, given the foundation from which Väyrynen derives his hedged principles, they seem not really to be doing any work. Väyrynen claims that when we have a reason not to lie there is some property of the situation that makes it the case that the reason not to lie is the reason it is - this is the 'designated normative basis' of the wrongness of lying. So, whenever it is wrong to lie, it is because the designated normative basis of the wrongness of lying is instantiated. But once we have a handle on what the designated normative basis is, it is unclear why we need to consider a principle about lying thereafter. We have all we need to know - some property bearing a relation to wrongness - which may or may not be distinct from lying. It is unclear how a principle articulating

the cross-contextual wrongness of actions instantiating the designated normative basis does us any good apart from the *in situ* apprehension of the designated normative basis. That is, it is unclear why other situations' moral properties have anything to do with the properties instantiated here and now.

But if the preceding considerations are roughly correct, then, we ought to step back and take stock of the landscape. If there is something in the skepticism expressed by particularism, then we ought to be wary of generalism. But the particularist views put forward by Dancy and McNaughton and Rawling have flaws or deficiencies that need remedying. Until that happens it does not look like we have a candidate for a strong particularist view. We can read the 'strong' here in one of at least two ways. First, a 'strong' view will withstand criticism. Second, a 'strong' view is one that is further toward one of the poles on a spectrum. If the work done and recounted in the preceding chapters is roughly correct, then extant particularist views are 'strong' in neither sense. Of course, we must be careful here not to set the standard for 'strong' in the first sense too high. That a view cannot answer a sufficiently radical skeptic might not be a mark against it. I myself happen to be rather convinced by the arguments outlined and presented in the preceding chapters – but that should not be a surprise. In conversation and correspondence, McNaughton has expressed his own skepticism at the arguments raised by McKeever and Ridge. That said, though, I have yet to see an adequate response to the objections raised by both Berker and Gert.[2]

At the same time, though, we ought to reject the two views that offer the best option for those with particularist sympathies who accept the untenability of Dancy's and McNaughton and Rawling's view - Lance and Little's defeasible generalizations and Väyrynen's hedged principles. But then that leaves us either giving up on particularism or accepting some variety of generalism. In this chapter, I will explore hope for a third option - that a revision to the basic particularist framework can provide a variety of particularism that is not subject to the same concerns that weigh the previous versions down. For instance, the views I have so far discussed are all subject to concerns about their employment of the Generalized Weighing Framework. Some of those concerns include the incoherence of the concept of a 'reason' employed in the Generalized Weighing Framework when that concept of a reason is holistic. Given that the view I go on to discuss in this chapter eschews that framework, it is not subject to those concerns. Additionally, holists, e.g. all particularists,

[2]That is, with the exception of *perhaps* Lechler (2012). As I discuss in a footnote in Chapter Two, though, I do not think that Lechler's piece is very much worth discussing in this context.

have some difficulties accounting for the apparent atomism of certain considerations. Among these, justice and cruelty seem to me the most difficult. Given the fact that the view I shall present is not a holist view, it again is not subject to any of these concerns.

So, the view that I will present in this chapter is a kind of atomistic account of reasons. I do not think I can argue, for certain, that generalists cannot co-opt this account. However, as particularists have done before me, I here register my skepticism that such a move could be successfully pulled off. In the coming pages, I hope to illustrate why. However, even if the view that I describe fails for some reason, it remains the case that extant particularist views are in need of revision. Even if there are defects in the current view, and even if they are significant ones, we will find the answer, if there is one, by looking more toward the particularist side of the spectrum of views than the generalist side. This confidence is justified, again, by the fact that the views I have discussed so far are problematic in the various ways that I have discussed.

In this chapter, I sketch a view that I am calling 'Eliminativism'. I call it 'Eliminativism' for one, or perhaps a couple of related, reasons. In short, the view is eliminativist because it rejects the traditional theoretical paradigm. It rejects some aspect or combination of aspects of the traditional theoretical framework. In the present case, it is an attempted elimination of contributory reasons talk from the discourse on particularism and reasons. In the first place, working within the contributory framework has only taken particularists so far and it does not look like it can take them any further. In fact, it is the conjunction of holistic contributory reasons and the Generalized Weighing Framework (GWF) that seems to lead to the particularist's most severe problems. At least, according to Berker (2007) and Gert (2007), it is this conjunction that renders the particularist's, in particular Dancy's, conception of a practical reason incoherent. And so I also reject the GWF. As I have said before, I believe that the GWF is artificial and overly intellectualized. Whatever standards we have for theoretical adequacy, to the extent our theories have empirical implications or rely on empirical matters, the plausibility of the theory bears a direct relation to those empirical matters and facts. In short, if the GWF depends on an objective weighting of reasons, then it implies a kind of structure to the weights of reasons that is empirically implausible. We do not weigh our reasons in the way implied by the GWF. We do not weigh a reason of strength value 6 against one of strength value 2, for instance. The imposition of strength value numerals on reasons, which seems necessary to make sense of the GWF, is entirely

artificial. It is a philosophical artifice. "But," the contributory theorist interjects, "these numerals are just a representation - they are not meant *actually* to refer to the strength value of the reason. Additionally, the GWF needs only relative weights of reasons to make sense, it does not need anything so precise as numeric weights for reasons. It would be sufficient for us to know that A is a stronger reason than B and B stronger than C for us to do our reasons calculus."

First, I can accept that the proponent of the GWF only needs a representation and does not need to give such a precise account of reasons as that one is of strength value 6 and another 2. But even if the GWF were to be amenable to, for instance, an ordinal scale – a scale according to which certain reasons were ranked higher or lower than others – this would not, I think, solve the problem. Suppose we have seven reasons, A through G, in a situation. Knowing that, for instance, reason A is the highest ranked reason does not help us to determine what to do when A counts for Φ-ing while reasons B through G count against Φ-ing. The question immediately becomes how much higher on the scale is A ranked over the other reasons. But without quantifiable values, we run into problems comparing reasons against one another. So the fact remains that, even in what we might think is the paradigmatic case of weighing reasons – say writing a 'pro' and 'con' list – there seems to me nothing going on that is so precise as the assessment of *individual* considerations. Rather, it is the entire list – the 'pro' side versus the 'con' side that gets weighed. Now, on the other hand, if the GWF permits the inclusion of subjective reasons (and it probably should) then, as I hope will be clear from the discussion to follow, it is in no better position than the view I go on to discuss.[3] Indeed, because of the problems raised by Gert and Berker, the traditional view might even be in a *worse* position than Eliminativism. And that is reason enough for us to begin looking beyond contributory reasons. To be clear, I am not at this point suggesting that the contributory theory is entirely false or incomprehensible. My point is just that the contributory theory seems not to be the best framework upon which to build a particularly *particularist* position. For that reason, I aim to look for another option.

Now, before I go on to describe in more detail what I take Eliminativism to be, I should make a few remarks about what kind of view I am offering. I take the view I am calling 'Eliminativism' to be one of a potentially varied class of views which, insofar as they reject the GWF, holism, and

[3]This distinction that I am employing between 'objective' and 'subjective' is just the familiar one that we find in Williams (1979)'s famous example of the agent who has mistaken gin for gasoline. Given the agent's error, objectively, she has no reason to drink what is in the glass – it is gasoline. Subjectively, having mistaken the gasoline for gin, she has a reason to drink what is in the glass – she believes it is gin.

contributory reasons, are appropriately called 'Eliminativist' views. As I am envisioning matters, these could be a wide variety of views employing somewhat, or even radically, different theoretical apparatuses. For instance, I will go on to discuss what I take to be an agent's reason for action. One element that I will describe as relevant to what reason an agent has to Φ is the circumstance in which the agent finds herself – taken as a whole, the entire state of affairs. I do not, however, think that this is a *necessary* component of all Eliminativisms. Another potential Eliminativism might take only subsets of the circumstance to comprise either a greater or lesser part of the agent's reasons, or perhaps even no part of the agent's reason at all. As long as the view rejects some component aspect of the traditional framework – either holism, contributory reasons, or the GWF – it counts as an Eliminativist view as I am using the term. Therefore, any theorist who is an atomist about reasons is, to a certain extent, an Eliminativist. Going forward I will discuss only a single brand of Eliminativism. For the sake of clarity and to avoid ambiguities, the discussion might even proceed as though the Eliminativism I am discussing is the only kind available. I will ask the reader to keep in mind, however, the fact that there are a variety of extant as well as potential Eliminativist views, as I am using that term.

My Eliminativism is a view according to which what agents ought to do is determined, not by adding together the weights of various considerations to yield an overall case-for some course of action, and weighing that case-for, against the case-against that action (or against the case-for some other action). Rather, according to Eliminativism, what an agent ought to do is determined in a far more 'wholistic' manner. It is a function of who and where the agent is and wants to be *in conjunction with the moral truths*. Like all other particularists and generalists - like everyone at the table - I think there are moral facts and truths that are accessible to us. Though, again, like all other moral particularists and generalists, I do not have a very good story about how it is that we access these truths and facts. I do not take this as a mark against Eliminativism. At least, it is not a mark against Eliminativism that is unique to Eliminativism. The nature of normative force is, as I see it, as notoriously mysterious as any other problem in all of Philosophical inquiry.

And so, in this chapter, I will explore the prospects for an elimination of holistic contributory reasons within the GWF from the particularist framework. The view I will advance is a significant revision to the basic particularist framework. I will argue that it can provide a brand of particularism that is not subject to the same concerns that weigh other extant versions down. First, since

the view I will advance is not, strictly speaking, a holist view, it is not subject to the concern raised by McKeever and Ridge, viz. that holism can be co-opted by generalism, thereby removing from particularism its main motivation. Instead, the main motivation for the present view is that, as I have discussed, extant particularist views are problematic. Despite these problems, I still have particularist intuitions. In particular, I worry that generalist principles are doomed to fail. I also worry, as I discussed in Chapter 4, that the deployment of defeasible generalizations or hedged principles requires that agents have the kind of knowledge and know-how that the generalizations and principles are themselves supposed to provide for the agent. If these concerns are warranted, then the search for a new brand of particularism is justified.

5.2 Beyond Contributory Reasons and Toward Eliminativism

As I have discussed at length, many current theorists conceive of practical reasons as contributory reasons. Such theorists conceive of contributory reasons as facts that (somehow) favor some course of action, Φ. That is, S might stand in some relation to a fact R in C such that R counts in favor of S's Φ-ing. We say that R is a contributory reason because it is a feature of the situation which contributes to (i.e. normatively favors) the case for S's Φ-ing. R need not be the only reason in the situation which favors S's Φ-ing and it might combine with other reasons for S in C. These other reasons might themselves favor Φ-ing or they might count against Φ-ing. For instance, the case for S's Φ-ing might be comprised of reasons R, S, T, L, N, E. The case against S's Φ-ing, or the case for S's Ψ-ing instead, might then be comprised of reasons F, G, H, and I. Whether S ought to Φ or Ψ depends, roughly, on the weights of the reasons that favor these different courses of action. I have mentioned that the Eliminativist's framework is rather different. In the rest of this chapter, I will explain, as much as one can, what the Eliminativist position is, the ways in which it is not subject to the same concerns as the traditional contributory framework, and I will attempt to respond to a couple of objections I anticipate traditional theorists might have against Eliminativism.

For the Eliminativist, reasons are neither holistic nor contributory, nor are they individuated in the way the contributory theorist individuates them. Nor is the Eliminativist an advocate of the GWF. Reasons do not weigh up, for the Eliminativist, the way they do for the contributory theorist. In fact, there is a real sense in which reasons do not 'weigh up' at all. Allow me to explain.

According to Eliminativism, context really is everything. And despite holistic particularist claims that they are privileging context, the theoretical devices they employ to explicate their theories seem to contradict this claim. If context really were *everything*, then particularists would have no need of default reasons or any device that is, effectively, a stand-in for a moral principle. After all, a principle is in a sense something in addition to the context, since the principle exists, as it were, outside of the context. If context were everything and particularism were true, then, ceteris paribus, we would need just to be in a context to figure out what we ought to do.[4] However, if indeed context is everything, then it will be exceedingly difficult to say much of anything about how it is that we figure out what morality requires. Eliminativism is in part an effort to make some kind of sense of how it is that we can, do, and should make justified moral decisions without appealing to the kinds of individuated considerations that get particularists into the kind of trouble already discussed. But Eliminativism can only offer so much. If indeed context is everything, then, again, it follows that there is actually very little that we can say from the outside.

Traditional particularists have been, to my mind, too ready to accept the terms of the debate set by generalists. Generalists have carried on the debate as though we can do all of our moral thinking from our armchairs; as though theorizing is a substitute for the lived experience of moral life. This privileging of theory over practice seems to be a bedrock assumption of the generalist position. If only we can just figure out the correct moral principle(s), the thought seems to be, then we have at least most of the moral work done before we have ever found ourselves in a moral context, faced with a moral decision. It is very easy to say that I should or should not flip the switch, push the large person, shoot the native, or detach the violinist. But it is a different matter when we have the switch, the large person, the native, or the violinist right in front of us. It is this difference, between the armchair perspective we take as theorists and the actual lived experience of our moral lives that Eliminativism is trying to bring to the forefront of the debate. However, at the same time, as I have said, Eliminativists think that there is in fact very little we can say about what we ought to do in a particular situation if we are not in that situation. The details matter. But, 'very little' and 'nothing' are two different things. And so, what we have are a few related questions. What is it for R to be a reason for S to Φ in circumstance C? How should (and

[4]Of course, we would also need the standard faculties any moral agent needs, but I take that as a given – as part of what it is to be in a particularly *moral* context. Without these faculties, we could not properly say that one is in a moral context. Owing to its probable lack of moral agency, a tardigrade, for instance, never finds itself in a moral situation.

how does) S choose between Φ-ing and Ψ-ing in C, if not by weighing considerations against one another?

According to the variety of Eliminativism I am advancing here, reasons for an agent S to Φ in some circumstance C are normatively significant relations, R_1, R_2...R_n, holding between S, C, and the options available to S, O_1, O_2...O_n. We might think of the reasons as properties of the state of affairs, the circumstance, C, in which S finds herself. Each reason is a product of the complex of entities – the agent, her current circumstance, and an available option. We can represent each of the agent's reasons by making the triadic relation explicit. Suppose S has only two available options. S, then, will have two reasons: R_1(S,C,O_1), say, to Φ, and R_2(S,C,O_2) either to Ψ or to not-Φ. According to the traditional, contributory view, the deliberative process by which the agent chooses between her options is one in which the agent 'weighs up' her reasons to determine which is most weighty. But as the Eliminativist sees things, no such 'weighing' occurs. Instead, S, in choosing, evinces which option she believes she ought to choose.[5] That is, S's choosing shows which option *from S's own perspective* bears the most normative significance, i.e. which reason's normative force S most strongly feels. And that is the option that, given her perspective, she ought to choose. Which option's normative significance S most strongly feels can be a function of many different things. S's conception of morality can influence which option S conceives as most normatively significant, as can S's conception of who she is and what kind of person she wants to be. So, for instance, the fact that S, were she only to consider her current desires, beliefs, preferences, etc. sees R_n(S,C,O_n) as most normatively significant, does not mean that S straightaway acts so as to bring about O_n. Instead, S might reflect on the kind of person she wants to be. This gives agents room to improve themselves, morally; to reflect on the reasons they 'see' and the reasons they want to 'see.'

Grounding the justification of the agent's choice in her conception of the normative significance of the relationship between her situation and the available options does *not* mean that the agent is *morally* justified no matter what she chooses. I am not saying that the agent's choice to be a torturer of children is morally justified if that is what she most wants to become or if that is the option that most strongly pulls on her. Though, the agent might in fact be rationally justified,

[5]The temptation is strong here to use the phrase, "...which option S *has most reason* to choose." The Eliminativist cannot, I think, allow such a locution. It seems to imply a kind of combining of reasons that does not happen within the Eliminativist framework

given means-end rationality.[6] Again, like all other theorists engaged in the particularism/generalism debate, I believe that moral facts exist and that there is a fact of the matter about which option is the morally correct one. The agent might be epistemically justified for choosing, say, O_1, given what she knows of her current situation and what she can reasonably expect from choosing that option. But whether the agent is morally justified is a matter of whether or not she is apprehending and appreciating the relevant moral facts in the situation.

How do we determine whether the agent is apprehending and appreciating the relevant moral facts? This is a difficult question. Indeed, this is a question to which I have never seen a satisfactory answer. As I say above, the nature of normative force is one of the most notoriously mysterious problems in all of philosophical inquiry. And the nature of particularly *moral* normative force is no less mysterious. For some normative domains, rationality for instance, normativity is not so difficult to understand. Frequently we can appeal to prudential concerns to drive home the irrationality of certain actions or beliefs. But there seems not to be a relevant analog in the moral domain; nothing to which we can appeal to drive home the force of moral normativity for someone who does not see it. In short, it seems entirely opaque how to determine whether an agent is apprehending and appreciating the relevant moral facts in a particular situation. There is no checklist we can look at to see if the agent has apprehended matters aright. Of course, we want to say that if S most strongly feels the pull of O_n, the option to become a torturer of children, and so sets herself that task, she is doing something morally wrong.[7] We want to say that S has misapprehended the situation, that she is making a mistake. And it seems clear that, were we to come across a person who has set herself the task of becoming the world's best torturer of children, we would rightfully regard her as morally depraved. But in keeping with the 'context is everything' mantra, according to Eliminativism, there is little that we can say about *why* this is the case from outside of the state of affairs in which someone actually sets herself this task. It might be that we are correct that she would be doing something morally wrong by trying to become the world's best torturer of children. But the Eliminativist's commitment to the primacy of context and to the non-existence

[6] That is, the view that it is irrational not to take necessary means to one's chosen ends strikes me as very plausible. This is not a problem for Eliminativism, at least not uniquely. Everyone has to deal with this kind of problem.

[7] That is, assuming there is nothing about about the circumstance C that would make becoming a child torturer the preferable option. We might imagine an anti-Neverland world in which a group of well-armed children are plotting to unleash a biological or chemical weapon that would destroy all life on the planet. In such a case, it seems to me morality *might* require that some of us become child torturers in order to extract information from the child-terrorists. That, of course, is not what is happening in the present situation.

of moral principles means that without being able to appeal to the particular situation, we cannot offer an explanation as to why it would be wrong. And so, ultimately, there is a real sense in which there is little that we can say to justify this position that would not, ultimately, amount to table-thumping and repeating our own moral intuitions. I take it that, in a way, this is at bottom what is happening when many, though perhaps not all, theorists argue over which moral principle(s) are the correct one(s). And so, I am afraid that like everyone else, the Eliminativist's response is a rather unsatisfactory one. Moral facts determine, from the perspective of morality, what S ought to do. But what S ought to do, lest she behave subjectively irrational, is what her reason $R_n(S,C,O_n)$ in the situation as she perceives it 'tells' her to do. This means that Eliminativism has the unfortunate consequence, like many other theories, that sometimes it is (subjectively) irrational to do what one ought, morally, to do.[8] Though this is an unfortunate theoretical consequence, it is one that we should expect if what we want is a theory that reflects the complications of everyday life.

5.3 Objections and Problems

Taking on board the Eliminativist framework would be a significant departure from traditional conceptions of reasons for action and perhaps, though to a lesser extent, to morality as an enterprise. It certainly does represent a fairly significant change for the particularist. Then again, leaving the Ptolemaic System behind in favor of the Copernican System was also a big departure for scientists. Being a significant departure from current prevailing perspectives is not, by itself a barrier to the

[8]But what of what is *objectively* rational for the agent? This is more tricky. I do not have a fully worked-out account of objective rationality, but I take it to be tied in some way to objective facts about particular agents' well-being. For instance, an action that will cause an agent some significant harm without any compensating benefit is probably objectively irrational. Objective rationality is a matter of having the ends that it is best for the agent to have – for some value of 'best' that I am, admittedly, unsure how to specify. So subjective rationality is a matter of agents reasoning well, arriving at judgments that accord with the logic of the beliefs and desires (among other things) that they currently hold. Matters are complicated further when we consider the fact that perhaps the well-being in question need not be the well-being of the agent performing the action. Perhaps, though an agent acts so as to decrease (or even destroy) her own well-being, the benefit to someone else's well-being makes the action objectively rational. The important point here is just that, according to Eliminativism, it is much more difficult to determine what is objectively rational than what is subjectively rational. Determining what is objectively rational requires knowing various things about the consequences of our actions that we, frequently, cannot know from within a context. So, whereas on my view, we can only make determinations of the subjective rationality of a judgment or action from within a context, if we can make any determinations of the objective rationality of an action, it seems we must be outside the context – we must be in a better epistemic position than I think we normally find ourselves. I want to say that an agent who, for instance, engages in practices of self-mutilation (cutting, perhaps) is acting objectively irrationally, though she may be acting subjectively rationally.

truth. There is *some* truth in Eliminativism, even if not all of the view is true. In this section, I discuss a few potential problems with Eliminativism. First, does Eliminativism preclude the possibility of akratic action? Second, is Eliminativism subject to a Gert/Berker style objection? Finally, is Eliminativism just the contributory theory in disguise?

5.3.1 Eliminativism and Akratic Action

Earlier, I said, "S's choosing shows which option *from S's own perspective* bears the most normative significance, i.e. which reason's normative force S most strongly feels." Does this imply that agents never choose contrary to the reason whose normative force they most strongly feel? That is, does Eliminativism rule out the possibility of akratic or incontinent action?[9] We can understand akratic action in a number of ways, but it is essentially action performed by an agent who judges it best to Φ but Ψ's instead (where Ψing might just be not-Φ-ing).[10] Does Eliminativism imply that akratic action never occurs? According to Eliminativism as I have described it, it looks like whenever an agent chooses to act, she chooses the option whose normative force she most strongly feels. But then it is never the case that an agent judges it best to Φ, but nevertheless Ψs instead.

What can we say here in defense of Eliminativism? In fact, I think this is only a minor objection to the view. We can simply translate the above Eliminativist claims in terms of the agent's judgments, rather than her choices. Rather than the agent's choice reflecting the option whose normative significance she most strongly feels, it is the agent's *judgment* which reflects which

[9] For a timeless treatment of the subject, see Mele (1987), in which Mele argues for the possibility of what he calls strict akratic action. While at one time I did not believe in the existence of akrasia, and though the arguments Mele (1987) advances are very good ones, it was a conversation with Alison Jagger in 2007 that convinced me that akratic action is in fact possible. And so, if it turned out that akratic action were impossible given the truth of Eliminativism, this is no longer a consequence I would welcome.

[10] There are a number of ways to describe what an akratic action is. Contributory theorists might prefer the following formulation in terms of reasons (or something like it): an akratic action is one an agent performs despite judging that she has better, sufficient reason to perform some other action. Mele (1987: 7) describes what he calls a strict incontinent action in the following way:

> An action A is a *strict incontinent action* if and only if it is performed intentionally and freely and, at the time at which it is performed, its agent consciously holds a judgment to the effect that there is good and sufficient reason for his not performing an A at that time.

I do not think for my purposes that the precise formulation matters. The basic and essential idea is that an agent judges it better to do something other than what she does. In so acting, she acts subjectively irrational–irrational from her own perspective, given that she acts contrary to what she judges her reasons dictate.

option she conceives of as being most normatively significant. The agent then, if she *chooses* a different option than the one she judges best, is acting akratically.[11] To the extent the Eliminativist wants to leave open the possibility of akratic action, then, and unless talking of agents' choices is meant to be synonymous with talk of agents' judgments, this might be a necessary adjustment to the view–and one I could easily allow. However, one reason for speaking in terms of what the agent chooses rather than what she judges is that, for the purposes of third-personal evaluation – evaluating the actions of other agents – there is nothing we can say about what she *judges*. We can only see the things agents choose (to do). And so, I find it much more natural to speak of what the agent chooses. However, I believe readers can substitute 'judges' for 'chooses' where relevant without any other changes being required of the view.

5.3.2 Eliminativism and the Gert/Berker Style Objection

One potential concern for my Eliminativism might be that it is subject to an objection like that which Gert and Berker raised against traditional particularism. Recall, those independently developed objections charged particularism with employing an incoherent notion of a practical reason. The complaint was that, given particularism's holism and, therefore, its lack of anything like the stable strength of a reason, the concept of a reason does not make any sense. Gert drew an analogy which helpfully sums up the problem. A fishing line which bills itself as a six pound test line should not break until six pounds of pressure is applied to the line. The six pound test rating only means anything if the line reliably holds up to pressure less than or equal to six pounds. If the line breaks erratically, sometimes when three pounds of pressure is applied, sometimes when thirteen, then it does not mean anything to call it a six pound test fishing line. Gert's objection relies on the notion that when discussing the strengths of reasons (or, the strength of a weight-lifter, a fishing line, an electromagnetic field, etc.), one is implicitly committed to a set of conditionals– conditionals about what would happen in other circumstances. So, to call a fishing line "six pound test" is to commit oneself to a set of conditionals like the following: "If the amount of pressure applied to this line is less than or equal to six pounds, it will not break. If the amount of pressure applied to this line is greater than six pounds, it will break." That is, talking of the strength of the fishing line (or reason, or weight-lifter, or electromagnetic field, etc.) commits one to *general*

[11] Thanks to David McNaughton for suggesting this modification to the view.

claims about what the fishing line can do in other situations. And one is so committed even when one is talking about a particular reason in a particular context. (Gert 2007: 552-3)

How might this objection apply to my Eliminativism? The concept of a reason with which I am working, that a reason is a normatively significant relationship had by an entire state of affairs, held between the agent for whom it is a reason, the state of affairs in which she finds herself, and one of her options eschews talk of the strength *of the reason itself*. Rather, if there is discussion of the strength of any aspect of the framework, it is the strength perceived by the relevant agent of the significance of the relationship to which she is party. That is, the notion of strength relevant for Eliminativism is not the strength of the reason, which language Eliminativism eschews. The only thing resembling strength on the Eliminativist picture is the 'strength' of the normative force felt by the agent. But that is in large part a consequence of her own desires, preferences, beliefs, etc., i.e. who she is and wants to be. The strength *of a reason* is only relevant if our theory permits weighing reasons against one another. But on the Eliminativist picture, no such weighing occurs.

Does it matter that Eliminativism eschews talk of the strengths of reasons if it still permits talk of the strength perceived by the relevant agent of the significance of the relationship to which she is party? Is this just a concealed way of talking about the strength of reasons? Or, if it is not, does the Gert/Berker objection simply have to shift its aim and thereby apply equally well to Eliminativism? After all, Gert's claim is that any discussion of strength commits one to the sort of conditional mentioned above. I think the Eliminativist has an easy, straightforward response to these worries. Even if talking of the strength of the significance of the relationship between the agent, her circumstances, and her option that the agent perceives were a concealed way of talking about the strength of a reason (and it is not clear to me that it is), the relevant conditional to which such talk would commit one seems not to cause Eliminativism any problems. This is because the Eliminativist's position is that what is relevant to what the agent ought to do in a given case is inextricably tied to the entire state of affairs in which the agent finds herself, as well as aspects of herself–her own psychology, motivational set, etc. This, as I see it, would render relatively inert any general conditional we could extract from claims about her reasons. Any such conditional, framed in Eliminativist terms, would be so specific that it would almost be an abuse of language to call it 'general'.

5.3.3 Eliminativism: Contributory Reasons in Disguise?

Part of being the kind of moral agents that we are means being able to explain our actions to others. Frequently we have to answer 'why' questions. Why did you meet Alice for lunch? Why did you give Brian that glass of water? Why are you going to the airport? The contributory theorist has a story about what kind of answers we give to these questions. That is, we cite the consideration which most strongly (or most saliently) counts in favor of the action we undertook. But what kind of response would the Eliminativist give? And is this response artificial? That is, is the kind of response the Eliminativist would give the kind of response that we *actually* give? Or am I, because of my Eliminativism, committed to a formula of response that is consistent with the theory but inconsistent with actual lived moral experiences? We can call this objection, that Eliminativism is just the contributory theory in disguise, the Disguise Objection (DO). On the one hand, the Eliminativist can cite a consideration, such as the fact that a promise was made or that so-and-so is in pain. But such considerations look like the kinds of contributory considerations that the Eliminativist eschews. Call this the C-Response. On the other hand, the Eliminativist could say something much more complicated, viz. citing a reason $R_n(S,C,O_n)$, but this seems too coarse-grained to be consistent with what we normally say in everyday contexts. Citing such a reason would amount to saying something like, "Given [C] the position in which I found myself, [S] the kind of person I take myself to be and want to be, and [O_n] the available options, Φ-ing seemed to me the best option." But this kind of response seems not to be part of folk responses to moral 'why' questions. Call this the E-Response. We have here a dilemma for Eliminativism. In response to one of these moral 'why' questions, either the Eliminativist gives a response that is consistent with the contributory theory of reasons or she gives a response that is inconsistent with what we naturally say in our everyday lives. In other words, either Eliminativism is the contributory theory in fancy dress, or it is, like I claimed the GWF was, over-intellectualized – a philosophical artifice.

How might the Eliminativist respond? Neither the C-Response nor the E-Response seem terribly satisfactory at first blush. I think there are a couple of ways that the Eliminativist might respond, though I will only offer what I take to be the most straightforward response. It would not surprise me if generalists (or other particularists) thought that the response I will discuss is not convincing. That is probably to be expected. After all, Eliminativism is already quite a departure from the traditional views. We need not accept such a high standard, though. I think it would be enough

to offer a response that is A] consistent with Eliminativism and B] not obviously bunk. To require that a response *convince* one's objector seems to set the bar too high.

By way of a response to DO, then, I suggest that Eliminativism might follow in metaethical subjectivism's footsteps. That is, we should interpret our everyday responses to moral 'why' questions as shorthand for E-Responses. So, when asked, 'Why did you go pick Steve up at the airport?' and I respond, 'Because he needed a ride,' this response is shorthand for some Eliminativist reason, $R_n(S,C,O_n)$. Spelled out more fully, my response 'Because he needed a ride,' is shorthand for something like, 'Because Steve needed a ride and I'm the only one in town he knows well. Given that we have a long history, it wouldn't feel right if I were to leave him to fend for himself. I care about Steve and our relationship. It would seem to me an affront to my own self-image if I were to leave a friend stranded at the airport. Etc.' I do not think that suggesting that our everyday responses to moral 'why' questions are shorthand for longer, more fully specified responses is the least bit unreasonable. And even if it were, the contributory theorist would be in no position to raise an objection to that effect.

Even within the contributory theorist's framework, we still might read 'Because he needed a ride,' as shorthand. That is, at least for the particularist contributory theorist, the simple fact that one's friend needs a ride from the airport is not by itself a reason. At least according to Dancy's view, such a consideration is not a reason without other considerations also obtaining. That is, there must be some relevant enablers, there must not be any relevant disablers, perhaps there also need to be intensifiers to increase the strength of the reason to pick one's friend up so that it outweighs another reason. In short, even the contributory theorist has to accept that responses to moral 'why' questions are often shorthand for the reasons that are actually operative. So the simple fact that the Eliminativist claims that everyday responses to moral 'why' questions are shorthand for some longer specification of a reason cannot be a cudgel against Eliminativism, at least not in the hands of other contributory theorists.

My response to DO, then, is essentially a straight denial of it. Eliminativism is not the contributory theory in disguise. The fact that our ordinary responses to moral 'why' questions do not transparently fit the formula of Eliminativism's reasons is just a consequence of how our language works. We often speak in shorthand, saying as little as we need to in order to convey as much meaning as we need to get our point across. This sort of thing happens in non- moral contexts

all the time. 'Why did Stuart Pearce get fired from the Nottingham Forest job?' To which an adequate response is surely, 'Because he's not a good manager.' But this response surely does not tell the entire story. There are myriad things that it does or can mean to be 'not good manager'. Pearce is tactically naïve, he fails to set his players out in formations appropriate for the grounds in which they are playing and for the teams against which they play. He was unable to cope with the pressures of such a high-profile position. He failed to man-manage his players, such that they exhibited a lack of discipline on the pitch. The list can go on and on. But it is too verbose and uneconomical to spell out all of these details when answering the question. But just because 'He is not a good manager' is the response that we give, does not mean that these further specifications are not true.

5.4 Conclusion: A Final Worry Alleviated

In this chapter I attempted to offer an alternative account of practical reasons for particularists to adopt. I suggested that because of the variety of problems particularists have, not least of all with the GWF, sticking with the contributory theory of reasons is effectively granting too much ground to generalists. The task I set myself was to develop a framework that is not subject to the kinds of objections particularists have faced in the past. I do not think the framework I have developed could be co-opted by generalists for the same reason Eliminativism is immune from the Gert/Berker-style objection. Given the degree to which Eliminativism privileges context, and given the amount of information built into the Eliminativist's account of reasons, any alleged generalizations we could build from those reasons would, effectively, just be rules governing particular situations. That is, they would not really be generalizations at all.

There is, however, one problem that particularists face which I have not yet addressed in this chapter. That problem is how the Eliminativist can account for the apparent fact that certain contributory considerations seem to be invariant. That is, certain considerations always seem to count, morally, and they always seem to count in the same way. Considerations of justice are like this. How can it be wrong or bad to do what is just? Perhaps also considerations involving innocent pleasures are among these kinds of invariant considerations. What could be wrong or bad about innocent pleasure? This is one of the problems that holists face. How can holists maintain that all reasons' strength and polarity can vary across cases when certain considerations seem never to

do that. Dancy tries to account for this by suggesting first, that his holism concerns what *may* happen, not what must, and then by drawing a distinction between a reason and its content (Dancy 2004: 77).

Regarding the claim that holism involves what may happen, not what must, this seems to weaken the view so much that it would be surprising if were not true. That is, this modal maneuver seems to render the view trivial. Secondly, I think the alleged distinction between a reason and its content is a highly dubious one. What if not its content, is a reason? I cannot see any sensible answer to this question. Indeed, the question seems to me to be making a category mistake.[12]

Now while Eliminativism is not, strictly speaking, a holist view of reasons, I have, up to this point, not discussed the alleged problem of invariant considerations because I do not think Eliminativism can give a response. But this is only because Eliminativism does not take the kinds of considerations that are relevant for this problem – e.g. considerations of justice or innocent pleasure – to be genuine reasons. They are, again, shorthand for Eliminativist reasons of the form $R_n(S,C,O_n)$. Given Eliminativism's privileging of context, alleged considerations of justice only make sense *within a context*. And it is the context that carries all the weight, that does all the work. And this is Eliminativism's lesson. If we are going to be particularists, we must go all the way. Context is everything.

[12]No where does Dancy clearly give an example of the distinction he is trying to draw, between a reason and its content.

BIBLIOGRAPHY

Berker, Selim (2007), "Particular Reasons." *Ethics*, vol. 118 (1): pp. 109–139.

Blackburn, Simon (1981), "Rule-Following and Moral Realism." In *Wittgenstein: To Follow a Rule* (Holtzman, S. and Leich, Christopher M., eds.), pp. 163–87, Routledge.

Broome, John (2004), "Reasons." In *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz* (Wallace, R.J.; Pettit, Philip; Scheffler, Samuel; and Smith, Michael, eds.), pp. 28–55, Oxford University Press.

Dancy, Jonathan (1993), *Moral Reasons*. Cambridge, MA: Blackwell.

——— (2000), "The Particularist's Progress." In *Moral Particularism* (Hooker, Brad and Little, Margaret Olivia, eds.), Clarendon Press.

——— (2004), *Ethics Without Principles*. Oxford: Clarendon Press.

——— (2006), "What Do Reasons Do?" In *Metaethics After Moore* (Horgan, Terry and Timmons, Mark, eds.), pp. 39–59, Oxford University Press, (Originally published in Southern Journal of Philosophy vol. xli supplement (June 2003) pp.95-113).

——— (2007), "Defending the Right." *Journal of Moral Philosophy*, vol. 4 (1): pp. 85–98.

——— (2013), "Moral Particularism." http://plato.stanford.edu/entries/moral-particularism/.

Davidson, Donald (1963), "Actions, Reasons, and Causes." *Journal of Philosophy*, vol. 60 (23): pp. 685–700.

Eklund, Matti (2011), "What Are Thick Concepts?" *Canadian Journal of Philosophy*, vol. 41 (1): pp. 25–49.

Gert, Joshua (2007), "Normative Strength and the Balance of Reasons." *Philosophical Review*, vol. 16 (4): pp. 533–562.

——— (2008), "Putting Particularism in its Place." *Pacific Philosophical Quarterly*, vol. 89 (3): pp. 312–324.

Hooker, Brad and Little, Margaret Olivia (2000), *Moral Particularism*. Oxford University Press.

Joyce, Richard (2001), *The Myth of Morality*. Cambridge University Press.

Kearns, Stephen and Star, Daniel (2009), "Reasons as Evidence." In *Oxford Studies in Metaethics*, vol. 4, pp. 215–42, Oxford University Press.

Lance, Mark Norris and Little, Margaret Olivia (2004), "Defeasibility and the Normative Grasp of Context." *Erkenntnis*, vol. 61 (2/3): pp. 435–455.

——— (2006a), "Defending Moral Particularism." In *Contemporary Debates in Moral Theory*, vol. 6 (Dreier, James Lawrence, ed.), pp. 305–321, Blackwell.

——— (2006b), "Particularism and Antitheory." In *The Oxford Handbook of Ethical Theory* (Copp, David, ed.), pp. 567–594, Oxford University Press.

——— (2007), "Where the Laws Are." In *Oxford Studies in Metaethics*, vol. 2, pp. 149–171, Oxford University Press.

——— (2008), "From Particularism to Defeasibility in Ethics." In *Challenging Moral Particularism* (Lance, Mark Norris; Ptorč, Matjaž; and Strahovnik, Vojko, eds.), pp. 53–74, Routledge.

Lance, Mark Norris; Potrč, Matjaž; and Strahovnik, Vojko (2008), *Challenging Moral Particularism*. Routledge.

Lechler, Andrea (2012), "Do Particularists Have a Coherent Notion of a Reason for Action?" *Ethics*, vol. 122 (4): pp. 763–772.

Little, Margaret Olivia (2000), "Moral Generalities Revisited." In *Moral Particularism* (Hooker, Brad and Little, Margaret Olivia, eds.), pp. 276–304, Oxford University Press.

Mackie, J. L. (1977), *Ethics: Inventing Right and Wrong*. Penguin.

McGinn, Marie (1997), *Routledge Philosophy Guidebook to Wittgenstein and the Philosophical Investigations*. Routledge.

McKeever, Sean D. and Ridge, Michael R. (2006), *Principled Ethics: Generalism as a Regulative Ideal*. New York; Oxford: Oxford University Press.

McNaughton, David (1988), *Moral Vision: An Introduction to Ethics*. Blackwell Publishing.

McNaughton, David and Rawling, Piers (2000), "Unprincipled Ethics." In *Moral Particularism* (Hooker, Brad and Little, Margaret Olivia, eds.), Clarendon Press.

——— (2003), "Naturalism and Normativity." *Aristotelian Society Supplementary Volume*, vol. 77 (1): pp. 23–45.

——— (2009), "Benefits, Holism, and the Aggregation of Value." *Social Philosophy and Policy*, vol. 26 (1): pp. 354–374.

McNaugton, D. and Rawling, P. (2008), "Holism About Value." In *Challenging Moral Particularism* (Strahovnik, Vojko; Potrc, Matjaz; and Lance, Mark Norris, eds.), pp. 166–184, Routledge.

Mele, Alfred R. (1987), *Irrationality: An Essay on Akrasia, Self-Deception, and Self-Control*. Oxford University Press.

Moore, G. E. (1903), *Principia Ethica*. Dover Publications.

Raz, Joseph (1999), *Practical Reason and Norms*. Oxford ; New York: Oxford University Press.

——— (2003), *The Practice of Value*. New York: Oxford University Press.

Ross, W.D. and Stratton-Lake, P. (2002), *The Right and the Good*. British Moral Philosophers, Clarendon Press.

Schroeder, Mark (2011), "Holism, Weight, and Undercutting." *Noûs*, vol. 45 (2): pp. 328–344.

Väyrynen, Pekka (2006), "Moral Generalism: Enjoy in Moderation." *Ethics*, vol. 116 (4): pp. 707–741.

——— (2008), "Useable Moral Principles." In *Challenging Moral Particularism* (Lance, Mark Norris; Potrč, Matjaž; and Strahovnik, Vojko, eds.), pp. 75–106, New York: Routledge.

——— (2009), "A Theory of Hedged Moral Principles." In *Oxford Studies in Metaethics*, vol. 4 (Shafer-Landau, Russ, ed.), pp. 91–132, Oxford: Oxford University Press.

——— (2012), "Thick Concepts: Where's Evaluation?" In *Oxford Studies in Metaethics*, vol. 7 (Shafer-Landau, Russ, ed.), pp. 235–70, Oxford University Press.

——— (2013), "Grounding and Normative Explanation." *Aristotelian Society Supplementary Volume*, vol. 87 (1): pp. 155–178.

Williams, Bernard Arthur Owen (1979), "Internal and External Reasons." In *Moral Luck: Philosophical Papers, 1973-1980*, pp. 101–113, Cambriidge: Cambridge University Press.

——— (1981), *Moral Luck: Philosophical Papers, 1973-1980*. Cambridge; New York: Cambridge University Press.

Wittgenstein, L.; Vescio, B.; Ogden, C.K.; and Russell, B. (2003), *Tractatus Logico-Philosophicus*. B&N Library of Essential Reading, Barnes & Noble.

BIOGRAPHICAL SKETCH

Richard Cordero was born in Smithtown, New York on July 2, 1983, suffering multiple birth-defects including a cleft foot, deformed pelvis, ventral septal defect, and scoliosis. He graduated from the State University of New York: The University at Buffalo in the Fall of 2005 with a BA in Psychology and Philosophy. He received his MA in Philosophy from The Florida State University in 2009. He also received his PhD from The Florida State University, in 2015. His research areas are primarily in Ethical Theory, though he has interests in Epistemology, the Philosophy of Action, Environmental Philosophy, and the Philosophy of Sport. Outside of philosophy, his main interests include soccer and his dogs.