THE UNIVERSITY OF PARDUBICE

FACULTY OF ARTS AND PHILOSOPHY

MORAL PARTICULARISM AND THE NON-FLATNESS REQUIREMENT

PETER NICHOLAS TUCK

DOCTORAL THESIS

2025

I declare:

The thesis entitled “Moral Particularism and the Non-Flatness Requirement” is my own work. All literary sources and information that I used in the thesis are referenced in the bibliography.

I have been acquainted with the fact that my work is subject to the rights and obligations arising from Act No. 121/2000 Sb., On Copyright, on Rights Related to Copyright and on Amendments to Certain Acts (Copyright Act), as amended, especially with the fact that the University of Pardubice has the right to conclude a license agreement for the use of this thesis as a school work under Section 60, Subsection 1 of the Copyright Act, and that if this thesis is used by me or a license to use it is granted to another entity, the University of Pardubice is entitled to request a reasonable fee from me to cover the costs incurred for the creation of the work, depending on the circumstances up to their actual amount.

I acknowledge that in accordance with Section 47b of Act No. 111/1998 Sb., On Higher Education Institutions and on Amendments to Other Acts (Act on Higher Education Institutions), as amended, and the Directive of the University of Pardubice No. 7/2019 Rules for Submission, Publication and Layout of Theses, as amended, the thesis will be published through the Digital Library of the University of Pardubice.

In Pardubice on 28.05.2025

Peter Nicholas Tuck

Acknowledgements

This thesis is firstly for my mother and father, Anita and Nicholas, and my brothers, Richard, George, and Oliver, who always believed in me, and without whom, none of this would have been possible.

I would like to thank my supervisors, Anna Bergqvist and Niklas Forsberg, who encouraged me every step of the way. Anna was my first ethics teacher when I first attended university. She was the one who first introduced me to moral particularism, and suggested I should read and write about it. Since I left England and moved to the Czech Republic, Niklas has taken me under his wing and been ever supportive. As the most prolific reader of every version of this thesis over the past few years, he always challenged me to think about what I really want to say.

To Antony Fredriksson, Olli Lagerspetz, Ryan Manhire, Ondřej Beran, Laura Candiotto, Silvia Panizza, Tomáš Hejduk, Lesley Jamieson, Matej Cíbik, and Ondřej Krása, who all, in their own way, encouraged and inspired me in the writing of this work, and the ideas it contains. I would further like to thank Tomáš, Ondřej, and Ondřej for all the times they patiently helped me negotiate the Czech system, without which, I wouldn't be able to graduate.

To all the friends I made along the way, without whom, the experience would have been very different. Vladimir Lukic and Patrick Keenan, when you were here we always did everything together, read one another's work, and encouraged one another in our writing. An honourable mention has to go to Martin Mraz and David Rozen, who were here with us too, all the way.

To my cute girlfriend Viola Pražáková, who I met at the start of my final year, who has helped me with all my projects ever since, and been my biggest supporter. You helped me with everything from the posters for my origami class to the PowerPoints for my lectures, but most especially with the formatting of this thesis, since I don't understand how Microsoft Word works.

I would like to thank Olena Kushyna, who joined the programme on the same day as me, and with whom I adopted a very small, weak, and scared dog. So last, but certainly not least, this thesis is for the very small hound, Iris Murdogh, who I found and named not long after I arrived here, and the naming of whom I consider to be more of an intellectual and literary achievement than anything written in this thesis.

Abstract

In this thesis I begin by examining two common assumptions shared by realist metaethical philosophers. Firstly, that ethics requires some explicated, systematic ontological structure, and concomitantly, an epistemology that explains our skill in navigating this structure to actually make ethical decisions. E.g. on moral generalism, there is an identified list of moral principles, and agreed ways of finding out what the principles are, then applying them to decide what to do in any particular situation.

I then point out that moral particularism, which is the view that there are no moral principles, has been widely criticised for lacking such a structure, and therefore rendering the metaphysical moral landscape "flat". This is called the flattening objection, which leads to the apparently connected claim that particularists therefore cannot provide an epistemology: a way to make specific decisions. Without structure, there can be no related skill, and since they have no structure, they cannot explain skill.

On the first point, about structure, I point out that the flattening objection relies on the assumption that what is required is an understanding of the moral landscape as having a global shape. I argue that while it is impossible for us to have knowledge of this global shape, individual decisions, which are themselves always local, only need rely on local shape. And that since particularists actually have a clear concept of local shape, the objection is not fatal, as long as we can provide a clear answer as to how we make reliably correct decisions relying only on local shape.

On the second point, about skill, I say that particularists do have a clear epistemological picture of how to relate local structure to moral decision making: the weighing of reasons. I conclude that since their view apparently satisfies the need for local structure and local skill, and that since these are sufficient for successfully navigating moral life, we can discard the common metaethical assumption that global ontologies are required, and therefore, dissolve the flattening objection against moral particularism.

Keywords: Moral Particularism, Moral Generalism, Flattening Objection, Quietism, Moral Perception, Reason Weights

Anotace

V této práci začínám zkoumáním dvou běžných domněnek sdílených realistickými metatetickými filozofy. Za prvé, že etika vyžaduje nějakou explikovanou, systematickou ontologickou strukturu, a za druhé, epistemologii, která vysvětluje schopnost orientace v této struktuře, abychom činili etická rozhodnutí. Např. v případě morálního generalismu existuje identifikovaný seznam morálních principů a dohodnutých způsobů, jak zjistit, jaké principy to jsou. Ty jsou následně aplikovány při rozhodování, co dělat v konkrétních situacích.

Poté poukazuji na to, že morální partikularismus, tedy názor, že žádné morální principy neexistují, je hojně kritizován za to, že takovou strukturu postrádá, a proto činí metafyzickou morální krajinu "plochou". Tomu se říká zplošťující námitka, která vede ke zdánlivě souvisejícímu tvrzení, že partikularisté tedy nemohou poskytnout epistemologii: způsob, jak činit konkrétní rozhodnutí. Bez struktury nemůže existovat žádná související schopnost činit rozhodnutí, a také takovou schopnost proto nedokážou ani vysvětlit.

K prvnímu bodu, tedy ke struktuře, uvádím, že zplošťující námitka vychází z předpokladu, že je třeba chápat morální krajinu jako něco, co má globální tvar. Tvrdím, že ačkoli není možné, abychom měli znalost tohoto globálního tvaru, jednotlivá rozhodnutí, která jsou sama o sobě vždy lokální, se musí opírat pouze o lokální tvar. A protože partikularisté ve skutečnosti mají jasnou představu o lokálním tvaru, není tato zplošťující námitka pro partikularismus fatální kritikou, pokud právě dokážeme dát jasnou odpověď na to, jak činíme spolehlivě správná rozhodnutí, která se opírají pouze o lokální tvar.

K druhému bodu, týkajícímu se schopnosti, říkám, že partikularisté mají jasnou epistemologickou představu o tom, jak vztahovat lokální strukturu k morálnímu rozhodování: pomocí vážení důvodů. Jejich názor zjevně uspokojuje potřebu lokální struktury a lokální schopnosti činit rozhodnutí, což postačuje k úspěšné orientaci v morálním životě. Docházím tedy k závěru, že můžeme zavrhnout běžný metatetický předpoklad, že je zapotřebí globální ontologie, a tedy vyvrátit zplošťující námitku proti morálnímu partikularismu.

Klíčová Slova: Morální Partikularismus, Morální Generalismus, Zplošťující Námitka, Kvietismus, Morální Percepce, Vážené Důvody

Contents

Introduction

The philosophical questions discussed in this thesis occurred to me, largely in their final form, during my first encounter with moral philosophy, when we had to take the introduction to ethics class during my undergraduate studies. Moral particularism was on the curriculum, and after being introduced to particularism, all of the main insights particular to that view immediately seemed obvious and sensible to me.

I decided to keep reading about particularism, and quickly found that what had seemed obvious and sensible to me absolutely did not make sense to many philosophers. Quite the opposite, actually. And these differences in what can seem obvious or sensible, or not, to different philosophers, is in a way the central theme of this thesis.

Particularism is the metaethical view that there are no truly general moral principles, understood as rules of behaviour that hold for all people, in all places and times. Common examples of such principles are "do not tell lies" or "do not cause pain". In the first chapter of the thesis, I explain in much greater detail what this view is. After all, when I was first told about particularism as the view that denies the existence of such principles, I wondered whether it was therefore meant to be a form of metaethical antirealism, or even metaethical relativism.

Jonathan Dancy, to whom the particularist thesis belongs, has pointed out that his view sometimes engenders these confusions. Further, I immediately began to wonder how particularists answer a range of more specific moral questions, like how we learn right and wrong without principles, or how we explain our behaviour to others without referring to principles. And it turns out plenty of other philosophers have independently followed the same train of thought, too. All of these distinctions and questions are clarified in the first chapter.

In the second chapter of the thesis, I turn to examine how particularist philosophers have tried to meet the non-flatness requirement, and by doing so, also allay the critics concerns about those more specific phenomena, like education, or justification. I conclude that unfortunately all these responses fail, each in its own specific way, but also in a shared way.

In the third chapter of the thesis, I turn to explain exactly what this shared reason for failure is. And I argue that it is because the external moral reality that realist moral philosophers believe in is epistemically inaccessible, which means that it is impossible for anybody to have any complete, and therefore general or global, knowledge about it. While this is a negative claim, I go on to make a further positive claim: that this does not matter, because nobody needs

any such general or global knowledge anyway. All they need to do is decide about the situations they are actually in. And in the rest of the thesis, I consider more precisely how particularist philosophers ought to explain this.

In the fourth chapter, I discuss a popular way particularist philosophers have tried to explain our having moral knowledge without reference to generality or globality, which is as "case-based reasoning". Case-based reasoning is a model of reasoning where the current case is compared to other, similar cases we found ourselves in, or we know somebody else to have found themselves in. I argue that while this is surely *part* of the story, it doesn't explain why anybody has knowledge in the first place, before they experienced many cases to use as data for such comparisons; and it doesn't explain how they decide what to do in a situation which is entirely different to any they found themselves in before.

So, in the final chapter, I turn to consider what particularists can say basic knowledge is, and I divide this question in half. Firstly, we want to know why anybody thinks they have reason to do anything, and I conclude that this question actually has no particular answer, but instead we just accept that there are reasons for or against doing things. Secondly, we want to know how we compare or contrast these reasons to decide what there is most reason to do.

In the literature, there is a consensus that this ought to be construed in terms of reasons having weights; of being stronger or weaker. And I suggest that this stronger or weaker maps pretty well onto the "central" and "marginal" division provided by the non-flatness requirement, which was where we started anyway. I find that here, again, critics think particularists cannot provide a clear explanation of what it means for reasons to be stronger or weaker. But I provide an answer: that we know the weights of reasons by comparison. I think that this conclusion is enough to allay *all* the worries of all the philosophers I identified and discussed during the thesis, which, indeed, were largely the same worries that first occurred to me when I first looked into particularism myself.

1. What Responses to Particularism Have in Common, and What We Should Take from Them

In this chapter I proceed as follows. I explain some relevant points about the metaethical position called moral particularism, partly by way of explaining moral generalism, the position it defines itself against. I then examine a range of different criticisms of particularism and say what I think they all have in common. Firstly, they invoke the idea that the external, independent moral reality, as conceived by realist moral philosophers, both has and requires a certain general *structure*, whether or not this *general* structure is exactly a principled one. Secondly, that our ability to be morally competent, to live well and make good decisions, depends on our understanding that general structure.

The criticisms are all selected because they are they have a practical basis: they connect what we say in ethical theory with what we do in moral practice, and focus on what ethical or metaethical theories are meant to be *for*; what they are meant to explain. I point out that all of the selected criticisms have the same form: they claim that ordinary human moral life somehow requires or depends on the existence of general truths for us to live by, and they differ only in which moral phenomena they choose to use as examples to demonstrate this. I then turn to show how some moral philosophers have instead focused directly on this shared problem about the connection between structure and competence, and given it the general formulation of the "flattening objection". I finish by drawing out the salient features of this objection.

But before we look at those specific criticisms, what are they inspired by, and what are they criticisms of? They are inspired by moral generalism, and are criticisms of moral particularism. And moral generalism and moral particularism are the names of two rival, exclusionary metaethical theories (we might instead call them stances, or approaches). Accepting one implies rejecting the other, since they offer absolutely opposite answers to the same set of questions. I take generalism first, since particularism is largely defined in opposition to it. And generalism is the metaethical name for the otherwise familiar view that there are fixed and *general* moral and evaluative truths, which can be codified into *general* laws or rules, like *lying is wrong*, and *pain is bad.* These are moral principles, and this principled approach to ethics has always been the dominant one.

For historical peoples, and still the vast majority of modern peoples, ethical rules or laws are found in religion. Most of the major religions, which both historical and contemporary humans were and are overwhelmingly members of, include such a principled approach to

ethics. In Judaism and in Christianity we find the ten commandments, as well as many other general obligations and prohibitions. In Islam eating pigs is absolutely forbidden. In Hinduism, eating cows is absolutely forbidden. So, before we even consider such rules at the metaethical level, is it relevant to note that it is an empirical, anthropological truth at the level of *descriptive* ethics that existing ethical systems usually have a generalist shape.

In secular, *prescriptive* ethical theories promoted by moral philosophers, we find much the same thing. The two most popular ethical theories, consequentialism and deontology, have general imperatives at their core. Consequentialism's key insight is that goodness (however construed) is *always* right making, and badness *always* wrong making. Deontology's key insight is that certain actions are *always* obligatory, and certain actions *always* forbidden.

The point here is that the vast majority of contemporary metaethical and ethical theorising takes place within a paradigm that assumes the concept of the general ethical truth. Indeed, at the level of ethical theory, we can only have a theory like traditional consequentialism at all by allowing that ethical theory is the search for such a general truth; and moreover, an ultimate and basic one. We can only have a deontology proper if we again believe that there are general truths to be found in the first place, and think that ethical theory is simply a matter of identifying the correct general duties.

Particularism is the recent suggestion that there are in fact no such general truths, laws, rules, or principles, and that therefore theories constructed within such a paradigm will be fundamentally mistaken and unsound. If this is in fact true, apparently the vast majority of past and present ethical theorising will be based on a metaethical mistake.

The reader may expect at this time that I will explain the arguments *for* these conclusions, particularly for particularism, it being the newer and more surprising one. But the arguments for these views are quite orthogonal to my focus here. Suffice to say that a range of arguments have been proffered in support of particularism, and that in the literature there is no clear agreement on which one is most persuasive. My focus here instead is on problems that apparently arise from accepting the particularist conclusion itself, that there are no general moral principles, however we get to being interested in it. This thesis then, is for those who are sympathetic to the idea that general principles might not exist, and are interested in both the philosophical and practical ramifications that follow from such a thought.

I will instead, however, examine some arguments that aim to *indirectly* make generalism seem more attractive than particularism in regards to moral *practice*. I will elucidate the problem with claiming principles do not *exist* by way of considering what we do with the

principles we believe exist, what they are *for*. Again, the goal is not to consider whether particularism is true, but instead to consider the further problems that arise if we do think particularism is true; which questions nevertheless remain for us, that a generalist paradigm apparently allowed us to answer.

While these are basic descriptions of the two positions as usually understood, it is important to note that these are umbrella terms, insofar as generalism and particularism are meant to be complete, comprehensive metaethical theories. Metaethics is a field that covers, nonexhaustively, the more specific fields of moral ontology or metaphysics, moral epistemology, moral psychology, moral semantics or language, and moral practice.

In other words, in metaethics we enquire about the existence of moral truths, how we know those same moral truths either generally or specifically, the nature of moral thought and reflection, the meaning of moral statements and the relation of moral language to the moral reality, and the relation between these other things and everyday moral behaviour. So, from the headline claims of generalism and particularism (moral principles do or do not *exist*), the following more specific claims can be readily derived:

Generalism:

Ontological claim: “general moral principles exist”.

Epistemological claim: “the possibility of moral knowledge depends on general moral principles”.

Psychological claim: “the possibility of moral thought, contemplation, or reflection, depends on general moral principles”.

Semantic claim: “the meaning of moral statements depends on general moral principles”.

Pragmatic claim: “participation in ordinary, everyday moral life depends on general moral principles”.

Particularism:

Ontological claim: “general moral principles do not exist”.

Epistemological claim: “the possibility of moral knowledge does not depend on general moral principles”.

Psychological claim: “the possibility of moral thought, contemplation, or reflection, does not depend on general moral principles”.

Semantic claim: “the meaning of moral statements does not depend on general moral principles”.

Pragmatic claim: “participation in ordinary, everyday moral life does not depend on general moral principles”.

It is absolutely essential to note that both of these views are about moral *realism,* which is the idea that there are moral truths, and further that they are truths about the external world. They are mind independent, being grounded in the external world, rather than, in some way, in human minds. On moral realism, there is always a neutral fact of the matter about what is right or wrong, or good or bad, and these truths do not depend on anyone knowing them, or thinking them true. On realism, the claim that “killing is wrong” is held to be of a kind with “triangles have three sides”; the “truth makers” of these claims are external facts that do not depend on human belief or opinion. Like generalism, realism is the dominant paradigm that the majority of historical and contemporary ethical theorising takes place in, both religious and secular.

So we have two basic assumptions, only one of which particularism and generalism share: that moral truths are “real”, and that moral truths are general. Generalists endorse both, and particularists only the former. Particularists perfectly well agree that moral facts are real and independent, they just think their truth depends on, is grounded in, particular contexts, rather than standing outside any or all particular contexts. For the generalist, “lying is wrong” will simply always be true, regardless of the particular situation, because of the way the universe is. For the particularist, “lying is wrong” will only sometimes be true, but when it is true it will be in just the same way, because of the way the universe is.

A second point that will be essential for the reader to keep in mind is that the more specific versions of generalism listed above are positive claims: they tell us positively what *is* the case. Specifically, for our purposes here, the ontological claim tells us positively exactly what *exists*: principles exist. Particularism, on the other hand, restricts itself to negative claims which are more of an observation about how we ought *not* to talk about morality, leaving it open for there to be a range of other views about moral reality that are not principled. The point here is that while particularism *can* be cast positively as the advisory claim “do not use general moral principles to inform your behaviour” (and indeed, some philosophers, e.g. Gleeson, 2007, do write about it in this way) particularist philosophers have been mostly silent on how you *should* inform your behaviour. The insights of particularists are more modest and general:

to consider carefully what situation you are in, to think carefully about what to do, and so on. But these rather vague recommendations can be more specifically filled out in all sorts of ways.

Again, specifically on the ontological claim, particularisms headline thesis is that principles do *not* exist. The central claim of particularism, strictly speaking, is "not generalism". Its central claim has nothing to say about what *is* the case. While recently some philosophers have begun to make various attempts to positively fill particularism out in one way or another, and later in the thesis we will examine such glosses, particularism is a new topic and is still relatively open in this regard. Anyway, they are not essential to the focus of *this* chapter: our goal here is just to discuss the basic idea that there are no general truths, and what might be problematic about this starting point. If one surveys the traditional explanations of particularism in the literature, it is possible to glean (at most) the following, seemingly sparse picture of what *is*.

1. Any natural fact or property can potentially, in the right circumstances, have moral import. There is therefore no specifically ethical subject matter.
2. Any such fact or property having a certain import in one place, can have a different import, or no import, in another place.
3. Given the unlimited number of natural facts or properties, and combinations of such, no limited description can be given of how moral and natural facts go together. One situation will not tell us anything about another. Any such description would necessarily be unlimited.

Particularism has been found as controversial as it is apparently radical: particularists have been charged with misrepresenting the moral reality (e.g. Crisp, 2000, McKeever and Ridge, 2006), the definition of moral knowledge (e.g. Stangl, 2006), the nature of moral thought, the meaning of moral language (e.g. Whiting, 2007, Jackson, Pettit and Smith, 2000, McKeever and Ridge, 2006) and so on. Of course, none of these complaints are surprising, since it is a matter of both historical fact and of definition that accepting particularism involves rejecting the majority of past and present ethical thought.

But more specifically, these charges only have the meaning or weight that they ought to have in this context *because* particularism is a form of realism. Remember, what we are calling generalism is the combination of the two ideas that ethical truths are both general and real, and all the above complaints can also be made in a similar way against any view which

rejects the second horn about reality, But if the point was actually that general truths do not exist because *no* objective, independent truths exist, we would have a quite different debate on our hands. Consider the widespread rejection of relativism among moral philosophers. Relativism denies *all* independent truths, and, for those (e.g. Enoch, 2011) who think that morality *requires* the concept of independent truths, it could be criticised as implausible or deficient in that regard.

Another useful example to illustrate the point of the contrast here is the error theory of J. L. Mackie, who famously held that all moral statements are false *because* independent moral facts do not exist (Mackie, 1977). So while such positions very well do exist, they are not the topic here. That is not the problem here at all. Particularists actually agree with all of their opponents that independent truths exist: the dispute is entirely about the precise nature of such truths. This is why they can be charged with, for example, misrepresenting the moral reality. If their position was that there is no independent moral reality, they could hardly be guilty of misrepresenting it; all of the philosophers I discuss in this chapter are moral realists who think the independent moral reality exists, and think they know what it is like, and this forms the basis of their critique of moral particularism.

If my position is that unicorns do not exist, I can hardly be accused of misrepresenting them because I claim their horns are golden in colour, not silver. There are entire schools of metaethical philosophy that hold positions on moral language such as, for example, that moral statements fail to refer to anything, that true statements need to refer, and are therefore all false (this again is the error theory of Mackie (1977), and is the claim about moral language that goes along with the ontological version mentioned in the previous paragraph). But particularists and generalists agree that moral statements do refer, and that they refer to independent facts. This distinction will be essential to keep in mind as we move through the discussion, and without it the problems and questions referred to will not be correctly understood.

So, that particularism is a form of realism is what makes it an interesting or distinct view at all. While historically generalism has always had such a realist character, and the idea that there are general ethical truths fits so naturally with the mind independent flavour of realism that the combination of generalism and realism inspires no surprise. Particularism, however, as the combination of contextualism and realism, is meant to be surprising and radical for precisely that reason.

This is because all the sorts of "antireal" views naturally incline towards contextualism, and find it unproblematic to fit contextual claims into their more general pictures. On

metaethical relativism, for example, since we already hold that ethical truths are grounded in particular societies in particular places and times, rather than in some shared, global moral reality, we get the insight that ethics is contextual for free. If particularism was simply the idea that ethical truths are contextual, there would be absolutely nothing new or different about it. The combination of contextualism and realism is what defines the particularist position, just as the combination of generality and realism defines the generalist one. I will now reconstruct and comment on a range of objections to particularism, made by such generalist realists. But before doing so, I will make it clear to what end I will do so. I have three claims I wish to make at this point.

Firstly, as a conceptual or definitional point, I wish to say that the ontological claims (moral principles do/do not exist!) are the most fundamental of the above options. Again, both are claims about realist ethics, and what all realist positions have in common is a central focus on what is independently or objectively *real*, as the name implies. While realist philosophers may ground their definition of reality in different ways (e.g., naturalists vs supernaturalists) they are united in defining their positions against *anti* realism, which has a corresponding fascination with what is *not* real.

For a realist what we have moral knowledge *about* or *of*, what our moral statements refer *to*, is the independent moral reality, and therefore what we have to say about the possibility of moral knowledge is presumably in accordance with our views on the same. The same point can be made about each of the categories above. On this view, we could suppose that all claims in support of generalism (at any of the above levels) both *must* and *do* reduce to "principles exist!", and vice versa. While the concept of an ultimate moral reality is an abstract one, the reason why we posit such an abstract entity is intensely pragmatic; it is about what an external moral reality is *for*. Realists promote their view ultimately for practical reasons, and this is a theme we will return to throughout the chapter, and throughout the thesis.

But this explanation of mine shouldn't carry the day alone. Instead, more strongly, I will try to show that the philosophers I consider also assume something like the above. If my opponents had views that in fact explicitly do not rely on the concept of existence, this point would become epiphenomenal. So secondly, I wish to say that all of the following examples seem to have an awareness, sometimes tacitly implicit and sometimes obviously explicit, of the above understanding of the practical point of realist thinking; why it *matters*. Specifically, they focus on the utility of *general* realist thinking. While the justification for such existing truths is philosophical, the "pay off", as it were, is practical.

Thirdly and more strongly still, I wish to suggest that some of the following *beg the question* against particularism by way of assuming generalism in their premises. This is an extension of the previous point. While generalism and particularism are united in existing within a realist paradigm, they disagree on whether we ought to think within a generalist paradigm. Therefore, a neutral investigation of the two views ought not to simply assume the truth of generalism. My interest here is in how we can criticise particularism outside of such a paradigm; to criticise it on some independent terms. So, while I shall allow that the practical concerns these critics raise are relevant and meaningful, I shall not allow that we should automatically construe them in general terms. But if I can demonstrate that some of the criticisms of particularism I survey contain "principles exist!" in their premises, this will be the clearest possible demonstration of how existence is actually the key concept being disputed, rather than some secondary concept.

The four apparently different practical topics I will discuss are moral education, the public role of morality in society, intrapersonal justification, and the making of excuses.

1.1 Moral Education

I take moral education first. Moral education is surely a central part of everyday moral life, and how we engage in it, or talk about it partially depends on, and relates to, what we really think about moral language and thought, the acquisition of moral knowledge, and ultimately, as is our focus here, what we think morality *is*. If we look again at our list from above:

> Ontological claim: "moral principles do not exist".
> Epistemological claim: "the possibility of moral knowledge does not depend on moral principles".
> Psychological claim: "the possibility of moral thought, contemplation, or reflection, does not depend on moral principles".
> Semantic claim: "the meaning of moral language does not depend on moral principles".
> Pragmatic claim: "participation in ordinary, everyday moral life does not depend on moral principles".

We can see better why any particular metaethical view ends up including all these parts, because of the subject matter that it is about. While moral education is surely a pragmatic topic,

our understanding of it is inexorably intertwined with these other kinds of considerations. If we want to have some comprehensive philosophical understanding of moral education, while we begin with the most basic question of what morality *is*, and therefore what the subject matter is, we want to go on to talk about how people know what is right, how they do or should think about their behaviour, what they should take moral words and statements to mean, and how they relate all of this to their ordinary day to day lives.

Anyway, what do we mean by moral education? One common meaning is the moral education of others: where a teacher, or someone taken to be a moral authority (who, on a generalist understanding, presumably knows the moral rules) imparts them to those others, so they can know them too. Often, the focus is on the moral education of children: the initiation of the new generation into the moral way of life. But presumably this can also be the moral education of adults. In the prisons, at least in the west, one of the traditionally stated aims is not just the punishment, but the rehabilitation, perhaps moral rehabilitation, of the offender.

However, moral education can also be the moral education of ourselves. The iterative process where we seek to identify and internalise moral truths for our own self improvement. In either case, the traditional general approach is stubbornly axiomatic: one first learns the relevant rules oneself, before imparting them to others who have not yet reached this point. If we hold that there are no moral rules, what are we to say about moral education?

Bradford Hooker, in his engagingly titled paper "Moral Particularism: Wrong and Bad" was the first to attempt an explanation of why particularism fails to explain something intuitively important about moral education, concluding "I think particularists won't be able in the end to give an adequate account of moral education." (Hooker, 2000, 15). His elucidation of this thought is as follows:

> General principles are certainly an important part of moral education. Particularists must admit this, since it would hardly be helpful to start off children's moral education by telling them that what matters morally depends always only on the particular circumstances. Of course children start off learning about morality by learning general principles. They are told not to do 'that'—where the 'that' means not merely <u>this act on this particular occasion</u>, but this <u>kind</u> of act. Parents typically teach their children by saying such things as 'lying is bad', not 'lying is bad on this particular occasion'. (Hooker, 2000, 18-19, his emphasis)

Here, Hooker's point that a focus on particular circumstances would hardly be "helpful" is a clear conclusion about the impracticality of particularism. But surely, as a self avowed realist, Hooker would agree that the principles ought to be correct (i.e., ought to exist and us to be correct in thinking such) for this to be the correct view of moral education. While the focus is on what would be easier for children, one thing we cannot do is read him as proposing some kind of "noble lie" in agreeing to take a principled approach to education whether we are sure of them or not, purely for practical reasons. This would be a sort of epistemic welfarism, or epistemic pragmatism, about moral language, which respectively mean that we should believe what pleases us, or what we find useful to believe, amounting to something like a Deweyan understanding of moral education, which is fundamentally distinct from the basic position Hooker (and his interlocutors in the particularism debate) share.

In realist moral philosophy the orthodox theory of language and truth is the correspondence theory: that independent facts in the external world are the "truth makers" of moral propositions, where moral claims "refer" to, and correspond with, that external reality In the modern analytic metaethical moral realist literature, this view was first clearly introduced and argued for by George Moore (Moore, 1910-11) who was the originator of what we now call nonnatural realism, and has remained a central tenet in that literature ever since (Finlay, 2007, 843, Huemer, 2005, 39-44). So for Hooker to actually be correct here, the general moral principles simply must actually exist; generalism must be the correct metaethical paradigm. And pointing out reasons why it would be useful for them to exist is not itself an argument for their existence; that has to be argued independently of such considerations.

Further, ironically, we would presumably be relying on principles (at the bare minimum, some kind of external evaluative *general* truths) to be able to say what is *good* or *bad* for children *generally*. Even if we are apparently agnostic, proposing a view based in some general truth about what is good for all children in all places and times is presumably helping ourselves to our potential conclusions in our premises, and therefore begging the question. Without going so far as to provide a developed particularist theory of education (since that is not actually the point here) one thing we can say is that, if particularism were true, the *correct* way to educate some particular child would itself be a contextual matter.

But if we don't simply assume that moral education *must be* the dissemination and inculcation of principles, then how would we want to describe it, leaving aside particularism for the moment? The most neutral description we can find is surely that we want children to gain the ability, or skill, to decide what is right and wrong. And if realism is also true, there

will be nothing in the description of this skill that is at odds with what we think ultimately exists; indeed, these two pictures must be in sympathy.

So, while we reject Hooker's insistence that moral education just must be about memorising rules, we have to agree with the spirit of his point. If particularism is to be a plausible theory, it should perhaps be filled out with some positive (not merely negative) picture of what actually exists, and have a corresponding explanation at how agents become skilled at regularly understanding the same. In conclusion, an "adequate account" in realist terms will apparently be one that assigns some kind of structure to the external moral reality, which not necessarily need be a generalist one, and education will be understood in terms of developing the skill to navigate this structure.

In any case, Hooker is clear that particularists *must* take his concern into account. This is a recurring theme in the particularist literature: strong intuitions about morality, or pre-existing moral practices, are held to be central pieces of evidence in both directions. David Bakhurst, in his exposition and discussion of Hooker's position, explicitly refers to what is really at stake, and really under dispute:

> Why does Hooker think that particularists are stymied by moral education? One reason might be that moral education as actually practiced includes teaching children principles. We tell them it is wrong to steal, lie, break promises, and so on. We encourage children to deliberate about problems by applying principles to them: 'Harry, that would be a lie and lying is wrong, so you shouldn't do it'. And we extol the virtues of being 'a person of principle': someone who knows the moral rules and sticks to them. Can we plausibly claim, as the particularist must, that such educational strategies are at odds with the true nature of morality? (Bakhurst, 2005, 266)

What the particularist must ultimately refer to, then, is the true nature of morality, which is what I said above is really what all these sorts of comments have to do anyway. On realist philosophy, correct practice will *always* accord with a correct understanding of reality. While the positions of Hooker and Bakhurst seem to suggest that this cuts both ways (if principles exist, we ought to use them; if we are committed to using principles, obviously they must exist) but philosophically speaking, we can only go one way here. We only get the practical "pay off" if we think principles in fact do exist. As I said about Hooker above, we cannot just will the

principles into existence. If we wish to be realist moral philosophers, we have to argue for them.

And that, of course, is the basis of the generalist/particularist debate. And questions about the true nature of morality are sorted out at a much more abstract level than practical questions about educating children. Again, and the point really cannot be stated strongly enough; educational strategies are simply not the right sort of *proof* for whether some abstract concept does or does not exist, and Bakhurst's rhetorical question is ultimately misdirected here; it is a general truth that some piece of everyday moral practice can be at odds with what we philosophically argue to be morally the case, and indeed, the entire field of moral philosophy is filled with claims and arguments that disagree with some naïve concern. What will be plausible is what will be argued for, and the critics concern is really that (at the time of their writing) particularists had not done this to their satisfaction. But we shall wait and see whether particularists actually have said, or can say, something plausible in response here.

Moreover, there is a great range of empirical, historical proof that considering how things are commonly done can be misguided, either in method, or in the content being in fact false. Education in every field evolves based on new information in any case, and for most of history, students of medicine, alchemy, or astrology, for example, were taught facts or laws that have turned out to be straightforwardly false. Our apparently justified belief about formal or abstract truths informs what we then go on to teach; not the other way around. The strongest reading of the views of Hooker and Bakhurst would be that the idea of teaching morality without using general principles is simply a nonstarter.

But firstly, we have ample reasons to believe that we can end up being mistaken about the correct way to teach any particular subject, and secondly, we have to allow particularists the chance to offer some alternative picture, that aims to meet the concerns that drive the worries these philosophers have about education here. And apparently, whatever this alternative picture is, it will have its basis in the "true nature of morality". Mark Lance and Margaret Little voice a similar worry about moral learning from a particularist viewpoint:

> The view seems to imply that there is no structure to moral theory at all. Moral understanding must be simply a matter of accumulating a series of one-off pieces of insight (x is good here; y is good there; do z next Thursday in Pittsburgh), a picture which makes it puzzling how morality could be learned…or even discussed. (Lance and Little, 2004, 436)

I find this interjection more interesting insofar as it does not simply claim that *principles* are what is required (it leaves open exactly what is needed) and therefore approaches the issue neutrally. They instead wonder whether morality needs some *structure* for its successful practice over many situations to be learnt. The obvious answer here is not immediately that we need principles, but that we need to be able to learn some kind of *skill* of analysing any particular situation we find ourselves in, and this is surely true. But specifically, that the external moral reality should have *some* structure for this skill to align with.

This version of the complaint leaves aside the question of whether principles exist, and unlike Hooker's version, does actually succeed in being independent of that existential question. I will return to this point in more detail later in the thesis, but at this point, the reader should simply bear in mind that this is what is actually important about these complaints; the general idea that our picture of what abstractly exists does indeed shape what we think we can say about moral *practice*, whether principles are in the picture or not.

1.2 Public Morality and Social Order

The second concept I will consider is the public nature of morality. Hooker also wonders about the repercussions of particularism being accepted as the correct shared metaethical theory at the group level; it being accepted by our culture or society as such:

> Imagine a society of particularists – by which I mean a society of people who don't believe or use general moral principles, even principles about what counts morally in favour of actions and what counts morally against. What would the consequences be of pervasive belief in and use of particularism?…the consequences would be terrible. If I'm right, then moral practice needs principles in order to avoid these terrible consequences. (Hooker, 2008, 26)

Again, I wish to suggest that this complaint simply begs the question. Rather than arguing in favour of general truths, it simply declares that any view which does not contain general truths would lead predictably to terrible consequences: itself a general evaluative claim! But let us do our best to consider the point further. The next thing we might say is that taken in its strongest form, Hooker's worry is simply a strawman. And this is because nowhere in the particularist literature, or indeed *any* moral literature, can we find the view that *social* norms or rules should

be discarded. Take the following two examples of social norms which presumably we want to preserve: speed limits, and the prohibition against dropping plastic on the ground. For both of these examples, we can think of a situation where we might break the rule actually *for* moral reasons, in a way that does not either problematise the use of keeping the rules, or endanger their status.

I might break a speed limit while rushing to take a seriously ill person to the hospital (indeed, one can find real cases where the police showed tolerance to a driver who acted precisely in this way). I might be holding a plastic bag in my hand when I see a person in serious need of help, and I drop the bag on the ground to rush forward and help them. No actual moral or social practice operates in a completely fixed way in the first place, but equally, I have found no philosopher who claims that we should abandon social norms because exceptions can be found. Indeed, all recent generalist metaethical theories also specifically allow for the possibility of exceptions, just as actual police officers and judges do. Even when perusing antireal theories that deny *any* independent moral truths, I find no philosopher concluding that social norms like those mentioned above should be done away with.

Consider two of the strongest versions of antireal metaethical theories: relativism and emotivism. On relativism, the social norms of different cultures and disagreements and discrepancies between the same are actually used as *proof* for the lack of general ethical truths (Harman, 1975), but it is of course part of relativism that people should follow the norms of the societies they live in. On emotivism, it is held that social norms are *just* the result of our shared positive or negative feelings towards the relevant behaviours (Ayer, 1936), but again, this is not meant to problematise the belief that we ought to respect or follow norms; it is just a metaethical view about where we ultimately get these norms from. My point here is that social norms are independent of the concept of independent, universal or general truth anyway. If Hooker wishes us to think that social norms can only be grounded in both real and general truths, no one really seems to think so, and there are plenty of other options available. So firstly, it is simply a strawman to claim that any moral theory that doesn't cleave to absolutely fixed rules is terrible. Secondly, none of them *really* do; real practice does not go in the way Hooker imagines at all.

Putting this point aside, I wish to turn to my main purpose here; to consider this example in terms of my suspicion that these criticisms ultimately relate to the question of whether ontological generalism actually is true, or more strongly, assume it is true. And here Hooker seems again to be using generalism to support what is supposed to be an argument for

generalism. This is because a general truth about terribleness is appealed to in evaluating a particularist society. In a separate work, Hooker phrases the problem in terms of what would be generally good or bad for us.

> One of the points of our having a morality is to increase the probability of our conforming with certain mutually beneficial practices.…a moral view does seem unattractive if widespread awareness of its widespread acceptance would have very bad net effects on human well-being. And that is just what I think is the case with particularism. (Hooker, 2000, 30)

Now, public morality is part of what is meant to be encoded in the laws. While morality and laws are not coextensive, Hooker's worry does seem intuitive when taken together with our usual ways of thinking about public morality. One can be agnostic about the actual of existence of general truths while still seeing the point that there is something, putatively ineliminable, that they are meant to be *for*. This is what I take to be the spirit of Hooker's position. Stephen Darwall considers Hooker's point about recognised, acknowledged and agreed mutual practices as being both an essential part of a morality, and part of what a morality is for. In the context of these "presuppositions", he says:

> These presuppositions can be met only if what we hold people accountable for, moral obligations, are accessible to all, hence to those held accountable, as shared public knowledge. And practices of accountability must be able to be public also. Interpersonal answerability is public in its nature; justifying oneself to others is possible at all only within a shared space of public reasons. (Darwall, 2013, 178)

I will not continue belabouring my point about whether this essential public facet of morality is bad news for particularism, or what it tells us about the ontological nature of the particularism/generalism debate. What I wish to do with Darwall instead is recognise the move here to a more specific focus of a particular feature of public morality: to our need to *justify* our behaviour to other members of our society, and to agree that it *is* essential to such justification that we have mutually recognised reasons. It is simply built into the nature of

justification that it will involve me offering reasons for why I acted so, and you recognising them as offerable reasons.

1.3 Intrapersonal Justification

To properly understand justification, we should consider it together with two other moral concepts that are integral to the definition of a reason: motivation and explanation. Consider something I have done, and my wish to defend my doing so to a friend. Firstly, I have the reason or reasons why I actually did the thing. These are the motivating reasons, and they need have no relation to an independent moral reality.

While the word "reason" is being used here, what is really meant is just why I took myself to have done the thing; there is no requirement for my explanation to be rational. Indeed, a request for a motivating reason may often be met with the report of an emotional state. "I did it because I was angry", is a perfectly meaningful and satisfying enough explanation of my motivation. If you asked me why I stole bread from a supermarket, and I say I did it because I was hungry and desperate, that is my motivating reason.

Secondly, I have the justifying reasons, which are the reasons I think the thing that I did was right. These are the ones I offer to my interlocutor to actually *justify* why it was *morally acceptable* for me to behave as I did. To continue our example above, if I say "and it is morally acceptable for people who are hungry and desperate to steal bread from supermarkets" I am now not just reporting on my personal psychology, but instead offering a defence of my behaviour by attempting to refer to independent, shared moral truths. I do not actually need to be independently correct for my justifying reason to stand as such; it just needs to take the putative form of a moral reason.

Thirdly, I have the reasons why the thing I did, as a matter of independent fact, was right or wrong. These are the explanatory reasons, and are what realist moral philosophers are usually referring to when they discuss "reasons" in the most substantial sense. On realism, explanatory reasons are mind independent facts about what really, truly is right or wrong. Whereas with justificatory reasons, what is actually at stake (at least, in terms of agreement) is whether myself and/or my interlocutor *believe* they are correct. Anyway, the key point for our purposes is that on generalism, moral principles are explanatory reasons, and therefore always actually correct to offer in justification. Mere belief engendered by rhetorical persuasion is irrelevant when the proper facts are public knowledge. On generalism, simply mentioning an

explanatory reason in my justificatory narrative is enough. On the example above, we can imagine a moral principle "it is acceptable for hungry people to steal food" that is widely known and endorsed in our society.

Presumably, I will be "right about being right" when my justifying reasons match up with the actual explanatory reasons, which are the ones I can ultimately appeal to. If my interlocutor is unconvinced by my contextual justification, the strongest thing left for me to say is that some x "is objectively right!". On generalism, of course, principles *are* these explanatory reasons, and are conveniently available to be read off when independent justification is called for. So, the specific problem with particularism here is that actually being correct that my justificatory reasons match the actual explanatory reasons is not enough: I have to persuade everyone else to agree with me, on a case by case basis. Presumably the generalist advantage is that I don't have to start again from the beginning each time. In our society, once a principle is firmly entrenched, I can presumably use it to quickly show anyone, at least from my own society, that my behaviour was justified. One off moral judgements don't seem to have the same status, and therefore the same advantage. Something like this understanding of justification seems to motivate Rebecca Stangl. Like with Hooker and Bakhurst, a plausible, deeply held intuition is presented about what our moral theories ought to take note of, together with a description of our usual moral *practice*, and a nod towards the concept of a general truth:

> When we judge that a particular action is morally right, we often cite a general moral principle in support of our claim. Perhaps we judge that the action in question will be the only one that will treat all parties equally, and we therefore justify the judgment by appeal to the principle that it is right to do what treats all parties equally. Similarly, when a dispute arises about what would be the right thing to do in a situation, we may expect that the parties to the dispute will proffer reasons for thinking their various judgments are in line with plausible moral principles. On one straightforward interpretation, these practices presuppose the truth of at least two claims. They assume both that there are substantive general moral principles and that the appeal to general moral truths can play a significant role in the justification of particular moral judgments. (Stangl, 2006, 201)

Now, there are two possible ways we can read Stangl's claim here that we can presuppose the existence of general moral principles. The first reading is simply that we can presuppose that *those who refer to general principles* in ordinary life take them to exist, and this seems obviously true. But the second reading is that these behaviours ought to lead *metaethical philosophers* to presuppose that they exist. And I find this specifically problematic. To say that these practices *presuppose* the existence of *general* moral principles, if understood in the second, philosophical, sense, both begs the question against particularism, and inappropriately uses pragmatic considerations as an actual argument for some abstract truth.

Firstly, we can say that at most, these practices require the existence of *independent* truths, without them having to be general. This is appropriately neutral between generalism and particularism, since on particularism, in any situation, there always are independent truths, and therefore, explanatory reasons, available. The particularist rejoinder here would be that it is simply irrelevant whether what correctly explains *here* explains somewhere completely different.

But the particularist would otherwise agree on the *realist* point here; that the existence of independent explanatory reasons is a product of realism. But whether these truths are ultimately general or not is simply a separate question from whether they exist in the necessary way to play the relevant role in explanation and therefore, justification *here.* And this question would not be answered by this form of pragmatic argument. What is the case is that such considerations prove that thinking about the nature of independent truth is *related* to, and made *relevant* by, pragmatic considerations. But a consideration making a question relevant is not the same as that consideration simply providing the answer to the question.

Little, like Stangl, thinks that particularisms sparse ontology poses a problem for justification, and ultimately, for explanation:

> The worry expresses a set of philosophical concerns that particularism undermines the status of morality. For instance, many charge that particularism is incompatible with the very notion of justification. Justification, it is urged, proceeds by subsumption under generality. If there are no codifiable generalities governing individual cases, then not only have we no method for finding answers, but nothing can function as a reason to be offered in defence of any conclusions we do reach. (Little, 2001B, 163)

Here, as in my earlier discussion of Little on the topic of moral education, I wish to draw a distinction between methods for finding answers, and the exact nature of those answers; in other words, the difference between ontology and epistemology. A careful reading of the quote above, at most, only leads us to agree that we need to identify the *method* or *skill* by which agents identify reasons in a particular context. While one method is offered as an example (subsumption under general truths which are commonly known) this is just one method. All particularists need to do here is have a rival explanation of how particular truths are identified, and such an explanation would have an epistemic form, not an ontological one. I take Little's reference to "a reason" here to mean an *explanatory* reason, as defined above. Explanatory reasons, by their putatively neutral and independent nature, are problematic for the particularist story in a way justifying or motivating reasons are not.

1.4 Special Pleading and the Making of Excuses

Hooker identifies a further problem, which is nevertheless related to the previous two: special pleading, or the making of excuses.

> Moral particularism as a decision procedure provides far too much scope for special pleading and rationalization of self-serving action. In the real world, particularism is unlikely to lead people regularly to make correct moral decisions. On the contrary, because I think that, with particularism as a decision procedure, people would persuade themselves that what they wanted to do was, in the particular circumstances, morally allowed, I think that people's use of particularism as a decision procedure would regularly have terrible consequences. (Hooker, 2008, 26)

Presumably when we engage in such special pleading, it is when we are seeking to justify our behaviour to ourselves or to others, and if such a practice became widespread, it would form part of the breakdown of "public" morality as described above. Anyway, here we again have references to what can only be independent standards of correctness and of terribleness. There are at least two ideas here, I suppose. The first is a strict interpretation, where if our society has a shared set of general principles, we simply cannot get away with making excuses. We will straightforwardly be blameworthy, whatever story we tell.

The second, milder, interpretation is one that also fits better with the current direction generalist moral philosophy is going in. Generalist philosophers have moved away from the idea of absolute rules, and instead prefer to discuss rules that allow for exceptions and excuses (McKeever and Ridge, 2006a). On this view, it will sometimes be okay to make an excuse, but it will precisely be the case that "the exception proves the rule". Those rare excuses that are weighty enough will leave the rule untouched, and not lead us to doubt its integrity. The problem with particularism, then, will not be that it allows excuses, but it allows them far, far too loosely.

1.5 What They All Have in Common

I hope I have done enough to demonstrate that what all these practical concerns have in common is either an indirect or background assumption about reality, or an explicit, direct link with the nature of reality. Even when read neutrally, in all cases, an ability to say something *beyond* the usual particularist schema about the moral reality is held to be intuitively essential, or ineliminable. While I will return to these examples later, at this point I will return to that usual definition, before considering those philosophers who have commented absolutely directly on this shared, common worry. Anyway, our definition was:

1. Any natural fact or property can potentially, in the right circumstances, have moral import. There is no specifically ethical subject matter.
2. Any such fact or property having a certain import in one place, can have a different import, or no import, in another place.
3. Given the unlimited number of natural facts or properties, and combinations of such, no limited description can be given of how moral and natural facts go together. One situation will not tell us anything about another. Any such description would necessarily be unlimited.

The apparent outcome of this is to leave us with absolutely nothing to say about what might be relevant or important outside of any individual case. And these four clusters of responses to particularism about four apparently different topics (education, society, justification, excuses) are all grounded in the same more basic, shared point: that the external ethical reality must have *some* kind of shape or structure beyond what particularism barely allows, and that this question

is one independent of how we specifically seek to answer the practical questions at hand, or whether particularists have good answers for them *limited to that particular topic*.

Take justification for example. Particularists usually endorse the narrative theory of justification. The narrative theory of justification, which was introduced into the recent analytic metaethical literature by Alasdair MacIntyre (1981), has a descriptive and prescriptive part. The descriptive part is that, when we are called upon to explain and justify our behaviour to others, we respond by telling the story of the situation we found ourselves in, what we identified as important in the situation, and the eventual decision we judged was the correct one to make. Like any story, this moral story therefore has a beginning, a middle, and an end.

The prescriptive part is that this is how we *should* explain ourselves, because on moral realism, there is a fact of the matter about which reasons were really present in the situation, and therefore which was really the correct decision to make. So on the narrative theory, to tell our story is to point out these features of the world, to invite our interlocutors to also see their independent correctness. Here is Dancy on the same:

> To justify one's choice is to give the reasons one sees for making it, and to give those reasons is just to lay out how one sees the situation, starting in the right place and going on to display the various salient features in the right way; to do this is to fill in the moral horizon. In giving reasons one is not arguing for one's way of seeing the situation. One is rather appealing to others to see it the way one sees it oneself and the appeal consists in laying out that way as persuasively as one can. The persuasiveness here is the persuasiveness of narrative: an internal coherence in the account that compels assent. We succeed in our aim when our story sounds right. Moral justification is therefore not subsumptive in nature, but narrative. (Dancy, 1993, 113)

But this theory can also perfectly well fit with a generalist story. While it seems clearly complementary for particularists, who focus so much on the details of particular situations, to be narrativists, I see no particular reason for narrativists to be particularists. After all, there is nothing in the above that prevents principles being part of the story. I could perfectly well tell a moral story where I acted in a principled way. So is whether we endorse the narrative theory of justification or not related to our metaphysical worry? I think not. The same point can be made about all of the examples, and this is why I have abstained from mounting a defence of

particularism in response to each specific criticism. I only mention the narrative theory of justification by way of demonstrating that the problem with all these examples is absolutely not that particularists have nothing to say about education, or society, or justification (they very much do) but that these answers are not seen by the critics as specifically solving what they take to be a specific problem with particularism.

And on the metaphysical worry, while I have rejected the stronger version of the point that some of the critics surveyed have, namely that ethics simply requires *principles*, a charitable reading of these responses ought to lead us to accept the weaker version; that it requires *something*. And the point I have tried to bring out here is that if we leave aside the specific insistence on *principles*, we find that the concepts of *structure* and relatedly, of *skill* come to the fore. My provisional conclusion about these examples can be summed up in the following three points:

1. Realists are united in thinking that there is an external moral reality, and that this external moral reality is in some way essential for explaining and justifying our everyday moral practice.
2. The philosophers I survey think that this external moral reality must have some kind of understandable structure and that we need to actually understand that structure to successfully be moral.
3. I think this worry is fair, but it shouldn't lead us to automatically assume a generalist paradigm for such a structure; only that *some* kind of answer seems to be required.

Accepting particularism therefore means taking this worry seriously, and considering what kind of other answers are available. I will now turn from these philosophers who only indirectly imply this general worry, to survey those who aim to describe particularisms general problem with structure, and therefore with competence, directly.

1.6 The General Flattening Objection to Moral Particularism

And in the literature, this problem with particularism's description of the structure of the external moral reality, and concomitantly, with our ability to gain and transmit moral knowledge and all that follows from this, has been described directly as the flattening objection,

which relies on the somewhat poetic metaphor of the "moral landscape". The moral landscape, like an actual landscape, is populated by mountains, trees, landmarks, and roads that represent the textured nature of the external moral reality, and specifically, act as guides, letting us orient ourselves to know where we are, and where we can go from here.

We can point these landmarks out to one another when discussing the lay of the land, and have a shared "map" of the territory. The problem with particularism then, is that its "moral landscape" is apparently just an empty, flat steppe. If *any* piece of information can matter, or not, on a case by case basis, we can never map the territory, never remember anything, and never agree with others. Looking around us apparently gives us no clue where we went before, where we are, or where we ought to go. Sean McKeever and Michael Ridge, who introduced the now widely used landscape metaphor into the discourse, and were also the first to describe the problem in terms of this landscape being "flat", explain their view as follows in their comprehensive book length treatment of the particularism/generalism debate, *Principled Ethics: Generalism as a Regulative Ideal*:

> One danger that has not been lost on particularists is that their extremely ecumenical view of moral reasons for action threatens implausibly to 'flatten the normative landscape'. After all, even if we think that in the right context… shoelace color can provide a reason for action (which, to be clear, we do not), there is an important difference between considerations of shoelace color and considerations of pain, pleasure, promising, and the like. (McKeever and Ridge, 2006, 47)

Hooker suggested that moral education relies on there being a "relatively small subset" (Hooker, 2000, 19) of morally important information, and this specific worry is channelled more generally by McKeever and Ridge here. It will be "implausible" if any piece of information, however apparently important or trivial, can actually be the opposite. This is a theme we will continually return to in this chapter: the inclination of moral philosophers in this debate to wish to divide information into *at* least two sets, or lists, where we have a list of information that is always important, that serves to explain if not the entirety, at least the majority of our moral lives, and another of information that is less relevant, or simply irrelevant. Anyway, this shoelace colour example was originally borrowed from Lance and Little, who wrote:

> Depending on which case the comparison is made to, any feature may assume moral significance, from shoelace colour to the day of the week: after all, against a rich enough story, there are cases in which the change from Tuesday to Wednesday makes all the difference. (Lance and Little, 2003, 291)

Note that this is simply a neutral description of what the particularists view is, with specific reference to #1 and #2 from the limited list of positive particularist claims I provided in the previous chapter, which was:

1. Any natural fact or property can potentially, in the right circumstances, have moral import. There is no specifically ethical subject matter.
2. Any such fact or property having a certain import in one place, can have a different import, or no import, in another place.
3. Given the unlimited number of natural facts or properties, and combinations of such, no limited description can be given of how moral and natural facts go together. One situation will not tell us anything about another. Any such description would necessarily be unlimited.

The example about the days of the week was itself borrowed from Roger Crisp, who referred to the problem in terms of *domination,* rather than flattening:

> If particularism is true, we should perhaps be a little surprised that certain considerations dominate our ethical thinking, whereas others do not. An extreme particularist can offer no account of why, for example, causing serious harm to people seems to matter so often, and why it's being Tuesday when the harm is caused never matters. (Crisp, 2000, 36)

Again, an example is given that apparently[1] appeals to our strongest moral intuitions, and it sets us to wondering whether particularisms ontology is too sparse: if all we can say is that *any*

[1] Something that the reader should keep in mind is that intuitions have become recently popular again in anglo-analytic moral philosophy, after previously being popular in the early 20th century. The apparently commonsensical intuitions of the philosophers *as people* are offered as an uncontroversial starting point, and test,

fact can be as relevant as another, how are we to understand why some seem so dominant, and some peripheral? Again, it is suggested that we naturally want to divide or filter information into at least two sorts. Out of everything we notice in a situation, we want to focus on a certain piece or on certain pieces of information as being the dominant ones.

Particularists do not disagree on this, exactly; on particularism in any *particular* situation, some pieces of information will of course be dominant. But for the particularist, this *could* turn out to be shoelace colour, or the fact that it is Tuesday. The generalist wants to have it that these lists can be prepared, kept, and maintained in advance, outside and before some particular situation arises. Interestingly, Crisp refers the position which I described above, and is meant to be *just* the standard formulation of particularism, as *extreme* particularism. I will return to this adjective later, since it is meaningful and bears further thinking about.

Lance and Little later strengthened their position regarding shoelace colour: While their first paper simply points out that on particularism, shoelace colour can matter, a later entry adds an evaluation, making it clear that this is actually the problem with it:

> The morally wise person, one might have thought, is someone who understands that there is a deep difference in moral status between infliction of pain and shoelace color. (Lance and Little, 2007, 16)

This makes explicit the second half of the point we are trying to draw out over this chapter and the preceding one: it is commonly thought in realist ethics that our *skill* in navigating real ethical decisions, our practical wisdom, depends on our correct understanding of the external moral reality. The morally wise person, on one way of understanding this, is one who understands the general moral status of certain truths.

This is the practical purport of the flattening objection: if the moral landscape is really flat, then it is questionable how anyone can get to be morally wise. In relation to the more

---

of ethical theory. The dominant paradigm is one where it is not the task of the moral philosopher to invent moral truths from nowhere, but instead to explain what is already going on. This is what is meant when arguments in modern moral philosophy are called "intuition pumps"; if either the inspiration for or conclusion of some argument seems absurd according to our intuitions, then there must be something wrong with it. The purpose of my own work here is not to mount a general challenge against the having and mentioning/asserting of intuitions, but we should keep in mind that we are always free to endorse or reject some offered intuition, and this entire debate around the flattening objection is very much built on the intuitions that the philosophers involved have had about what might be called "common sense" morality. I will return to this point at the end of the present chapter, but for now, suffice to say that whenever the word intuition appears it should be taken to be synonymous with "common sense", which is how it is used informally in ordinary speech. I make this point because, in formal philosophy, intuition can sometimes be used as a technical term with a specific, different meaning. For example, in formal epistemology, intuition can refer to a sort of *a priori* knowledge which is gained without any empirical experience.

specific examples of important moral phenomena discussed earlier in the chapter, if nobody can have practical wisdom, how can the same be taught? How can society function? The problem has also been described in terms of relevance and irrelevance:

> A putative problem for the moral particularist is that he or she fails to capture the normative relevance of certain considerations that they carry on their face, or the intuitive irrelevance of other considerations. (Thomas, 2007, 77)

For Thomas, the relevance is "intuitive"; the considerations carry their relevance "on their face". Again, it seems to be simply obvious to many metaethical philosophers that a particularist way of thinking is suspicious, that it simply doesn't accord with our ordinary ways of thinking or acting. It simply *must* be the case that there are at least two sorts of moral information. The division has finally been given in terms of "centrality":

> On particularism, thick moral properties have no more intrinsic moral significance than non-moral properties. It will, presumably, turn out that these properties are 'commonly more important' than some others (although particularism owes us an account of why), but that not only understates their force, it seems to mislocate their centrality. It is not just that it is helpful to look at them first because they often count; their counting is central to their being thick moral concepts. (McNaughton and Rawling, 2000, 272)

McNaughton and Rawling suppose that we should simply *presume* that there is a "limited subset" of moral information, as Hooker puts it, that is more important than the rest, and write as if we already know that it will be helpful to focus on this sort of information. And indeed, this is how ordinary moral thought, as well as philosophising about the same, usually goes on.

Their version is particularly helpful in that it tells us that particularism "owes us an account of why". "Owes" refers to the potential of flattening being a problem that any successful theory must take note of. "Account" here is of course a philosophical discussion of the problem; but what counts as an account is internal to specific discourses. I refer to this because later, I consider what sort of "account" would be successful in this debate; not just any kind will do. Finally, "why" here, when taken in context, presumably refers to a fundamental explanation of what exists, e.g. what is the specific ontological status of a central vs peripheral

truth. The demand here is that particularism not only tells us why some truths that are obviously central are central, but why they got like that in the first place. And the worry is that their way of looking at things allows *no* answer to the question, when one is owed.

Anyway, I think this second set of examples is enough now to make the point: these metaethical philosophers worry whether particularism can explain something clearly obvious about morality, or more strongly, are certain it doesn't. And unlike the first set of philosophers I considered earlier in the chapter, who confined themselves to pondering about some specific feature of ordinary life like education or social order, these philosophers directly mount a *general* challenge to particularism. Vojko Strahovnik, who has written a series of papers discussing particularism and flatness, concurs with McNaughton and Rawling on this move of turning a "worry" into a demand, something that a complete metaethical theory *owes* us, which he calls the "non-flatness requirement", and which he has already concluded particularism cannot accommodate:

> The moral non-flatness requirement says that any moral theory must somehow account for the fact that some considerations or features of acts are more central to morality than others…radical particularism cannot properly account for the thought that some considerations are more central than others. (Strahovnik, 2016, 67-68)

So this is the central, common problem that motivates all of the more specific worries I previously surveyed. I now wish to further discuss a few things that come out of the above. Why, precisely, is particularism extreme or radical? Should we concede that it renders the moral landscape "flat"? Further, I will explain how the moral landscape metaphor allows us to understand the difference between general and particular truths, by recasting them in terms of globality and locality, and what this might mean for a solution, or at least, a response.

1.7 Is Moral Particularism Extreme or Radical?

Strahovnik, similarly to Crisp's description of the view as extreme, brands a particularist view that doesn't meet the non-flatness requirement as "radical". I think Crisp and Strahovnik's choice of adjectives deserves further discussion. What is it that makes a philosophical position "extreme" or "radical"?

One way to understand it is that the position is radical according to itself. On this understanding, there will be a standard view called "particularism", and naturally radical particularism will be a more extreme version of this orthodox position. Such language is commonly used to describe variety in any kind of normative stance. We are familiar with phrases like "moderate vegans" and "radical vegans", or "moderate Muslims" and "radical Muslims". It seems obvious that the meaning here is that the vegan or Muslim is meant to be radical *for* a vegan or Muslim, at least partially; one way they are radical is that they are radical when compared to the majority of their group.

The second, related and extended definition would be that these positions are radical according to some external, shared normative standard. E.g., whether or not the Muslim is radical *for* a Muslim, he is also or instead radical according to the normal standards of liberal society. Or if the vegan is radical, he is radical according to the views and behaviour of the majority, who eat meat. I think users of the word radical usually aim to include at least one, but often both of these sorts of points when they use the word.

But on particularism, I don't think the first understanding is correct, since the view being discussed simply *is* normal particularism. There is no majority of particularists who hold some other, milder, view, and a smaller group of hardcore fanatics who promote an extreme version of the position. There are just particularists, who think general moral principles do not exist, so, instead, it seems that ordinary particularism must be what is "radical". So if it isn't radical in a way internal to particularism, according to its own standards, then it must be radical according to some external standards, and therefore only by the second sort of definition.

So the question becomes, the standards of which shared community is particularism radical by? It is probably fair to say that particularism is radical according to the shared standards of most human societies; since in most societies, there is widespread belief in general principles. But this is an observation in descriptive ethics. The claim of radicalness here is being made by metaethical philosophers, who aim to stand outside of particular societies and make abstract, prescriptive, and universal points. So I think it is also radical according to the norms that are popular among the community of metaethical philosophers, but in a way which is intertwined with the first. Ordinary moral life is what metaethical philosophers take themselves to be philosophising in regard to. It is their data and their inspiration (one would hope).

At this point we ought to digress, and consider what is meant by ethical and metaethical philosophy. Metaethical philosophy is where we consider the background assumptions and

standards that shape how we formulate or evaluate ethical theories. Metaethics is meant to be where we propose neutral, external standards about how ethical theories should be understood. So if we call an ethical theory radical, it is because we think it is radical *according to* some neutral, external standard which has broad support among metaethical philosophers.

On this understanding, to call particularism, which is a metaethical theory rather than an ethical one, radical, is a metametaethical claim. It is radical *according to* some shared, neutral understanding of what a metaethical theory ought to be like. Just as philosophers have shared assumptions for how they evaluate ethical theories, surely they have shared assumptions for how they evaluate metaethical theories. And while more examples can always be given, I think the selection I have surveyed here are enough to show how particularism seems unusual to metaethical philosophers.

So branding particularism as "extreme particularism" or "radical particularism" is not simply to name the view, or even describe it. Instead, it both an anthropological evaluation of the view, and both separately and further, an evaluation in relation to the non-flatness requirement. This requirement neatly sums up what would apparently be required in response to the non-flatness requirement. Any moral theory must account for non-flatness.

1.8 Is It Accurate to Say That Particularism Makes the Moral Landscape "Flat"?

Now we have dealt with the apparent radicality, there is a point that I feel is important to make in relation to the way the flattening objection is framed and named. The complaint that particularism renders the moral landscape entirely "flat" actually seems better suited to describing a nihilistic metaethic, one where there are actually no moral truths. Such an external moral landscape would therefore really be empty. But, as I have emphasised before, particularists are realists who think there *is* an external moral landscape, and there *are* external moral truths. They just think there are many more than a generalist ethic allows.

Particularists do not say the external moral landscape has nothing on it – actually, they say pretty much the opposite. They say it has an unlimited amount of things on it. Their complaint is that generalist views actually have too *few* features; that generalist philosophers are thinking far too simply and quickly when they proclaim to know what is obviously central to ethics. If we take the deontic theory of W. D. Ross (1930) as an example, Ross says that there are seven general ethical truths. The particularist rejoinder is that this is too few, too

simple, that there are far more than seven – the amount is completely open. The point here is that, on generalist moral philosophy, a moral view can be either monist or pluralist.

Consequentialism, for example, is usually understood as monist in that it has one supreme moral principle that explains all moral phenomena. Views like that of Ross, that instead allow for a range of ethical truths, are therefore called pluralist. But particularism formally allows for neither monism nor pluralism, since it doesn't allow that there are any particular amount of general truths at all. Particularists really stand outside of the monist/pluralist divide (Dancy, 1983, 542). While a negative understanding of particularism, which I understand the flattening objection to incorporate, would be that they have no truths, the positive rejoinder is that they have endless truths; far more than either the monist or pluralist.

The point here is that the flattening objection seems to set up a comparison between a moral landscape with some limited amount of things on it, and a moral landscape with nothing on it. Really, the comparison must be between a moral landscape with a limited amount of things on it, and an unlimited amount of things on it. Since the preferred vocabulary in the literature when referring to the moral landscape as having features is to describe it as "textured", it may seem as though this problem could really be called "the problem of unlimited texture". Since this metaphor is meant to be a visual one, we might point to the difference between the surface of a mirror, and the surface of a towel, with the second in fact being more apposite. Because the problem is standardly known as one about "flatness", I shall however continue to refer to it as such.

Now we have considered what it means for particularism to be extreme or radical, and how precisely it renders the moral landscape "flat", I wish to return to the general suggestion that has been drawn out from the philosophers I have surveyed: that there ought to be at least two sorts or kinds of moral information, where particularism only allows for one. Amelia Hicks is particularly clear about the need for more than one kind: "Moral particularism 'flattens the moral landscape'… particularism treats moral reasons of different kinds as if they were reasons of the same kind" (Hicks, 2014, 12).

1.9 Two Kinds of Moral Information

We have the following pairs of concepts: Central and marginal, dominant and impotent, relevant and irrelevant. Recall that on nonnatural realism, which generalism and particularism

are varieties of, we have the nonnatural moral fact(s) which is/are *grounded* in some natural fact(s). Looking again at our description of particularism, we can see that it only allows for one sort of natural/nonnatural pairing: a *possible* one, with *always* and *never* explicitly ruled out by the formulation.

Whereas on generalism, we would have either two or three kinds of information. Firstly, those where a certain natural fact *always* grounds a certain evaluation. If we endorse the principle "lying is always wrong" then the natural fact that someone purposely said something they knew not to be true would always ground the nonnatural property of wrongness. Secondly, those where a natural fact *never* grounds a moral evaluation. For example, in McKeever and Ridge's introduction to their view, they told us "… even if we think that in the right context… shoelace color can provide a reason for action (which, to be clear, we do not) …" (McKeever and Ridge, 2006, 47). Thirdly, we might posit the category of natural and nonnatural pairings that only sometimes appear: this sort is of course the kind favoured by particularists, and is acceptable to generalists, too. Anyway, all of these examples mentioned in this chapter and the previous are meant to lead us to see that we ought to have *at least* two kinds, and what the possibilities are. Returning to the detail of one of those examples may be helpful in seeing this point. Here again is Crisp's summary:

> If particularism is true, we should perhaps be a little surprised that certain considerations dominate our ethical thinking, whereas others do not. An extreme particularist can offer no account of why, for example, causing serious harm to people seems to matter so often, and why it's being Tuesday when the harm is caused never matters. (Crisp, 2000, 36)

Crisp contrasts what is dominant with what *never* matters. This is a conscious further step, insofar as particularisms standard category of mattering is *can sometimes* matter. As we have seen, McKeever and Ridge also deny the *sometimes* category. If one were to add a second category, as some of these objections seem to imply, it seems one would naturally end up with:

1. Some facts always matter.
2. The rest matter only sometimes.

But the generalism of these philosophers doesn't even allow for this more moderate position. I think that holding that something can *never* matter morally is the strongest version of the generalism that particularism is opposed to. This point allows us to cast the flattening objection in both a positive and negative way. The positive part of the flattening objection, the one I have spent the vast majority of my time here seeking to draw out, is that there are some special ethical truths that ought to have some central status in our moral lives. The negative, opposite version of the view is that there are some truths that simply have no relevance to our lives. I think this distinction is important since it leads us to wonder whether all these philosophers have precisely the same point, or whether we need to rule on which version we wish to focus on. So Crisp (and McKeever and Ridge, and Thomas, for example) seem to suggest:

1. Some facts often matter.
2. The rest do not matter at all.

This detail perhaps allows us to divide the philosophers we have surveyed into two kinds. Those who insist that the "sometimes" category should be rejected are those entirely committed to generalism, while those who are sympathetic to the "sometimes" category, but still cling to the "always" category, make some allowances to particularism while still thinking it somehow falls short. Hooker, again on education, also concludes that only some facts will ever matter, while seemingly allowing for the possibility that some facts can matter sometimes:

> Sometimes moral particularists say moral principles are really just indications of the way a fact can count morally. This seems an especially weak answer if conjoined with the thesis, which some particularists assert, that any fact can be morally relevant. True, children need to learn which facts can be relevant. But they understandably hope this is some (relatively small) subset of all possible facts. (Hooker, 2000, 19, his emphasis)

Presumably this "subset" would consist of what is central to, or dominant in, ethics. Even if we were to otherwise agree, as I suggest, that any fact can *possibly* be relevant, there ought to be some main ones which are particularly useful to focus on; possibility might come in degrees. In his discussion of public morality, Hooker also tells us:

> One of the things a shared commitment to morality needs to do is provide people with some assurance that others won't attack them, rob from them, break promises to them, or lie to them…knowing that others have certain firm moral dispositions can give us added assurance how they will behave…now if shared commitment to morality should, among other things, create settled expectations about how others will behave, how does particularism look? Imagine we knew of other people only that they were committed moral particularists…Would we have enough confidence that they'd virtually never attack us, rob from us, break their promises to us, etc.? (Hooker, 2000, 22)

What is helpful here is that Hooker gives us a short list of what I suppose he takes to be the central concerns, and it reminds us that on generalist views, an explicit list of an exact number of *always* moral truths is sometimes provided. To mention Ross again (since his is a paradigm example) Ross claims to have discovered that there are seven central moral truths. A survey of all similar examples seems to give us the following possible pairings:

1. *Always* and *never.*
2. *Always* and *sometimes.*
3. *Sometimes* and *never.*
4. *Often* and *rarely*.

Now, I would suggest that to avoid begging the question, and to allow particularists to respond fairly, some of these formulations are too strong. To suggest always and never as options from the start seems to firstly beg the question in favour of generalism (since any "always" pairing seems to meet the definition of a principle) and secondly to beg the question against particularism (since they take the opposite position of excluding "never" from their basic view).

If we simply concede to McKeever and Ridge that shoelace colour can *never* matter, or to Crisp that it being Tuesday today cannot possibly be morally relevant, we are already discarding something that was meant to be important about particularism: that we ought to always look hard at the situation we are in, since it can be surprising what will matter, what will be relevant. So, Since this thesis is about what comes out of being sympathetic to particularism, I will ignore the line of thought that leads to declaring facts to be *always* or *never*

relevant, while allowing for the possibility that we can still have some robust notion of centrality that might be better understood in terms of what is often and what is rare.

This suggestion of mine where we instead allow for degrees of possibility, and therefore for centrality and marginality, seems to preserve the spirit of the problem while avoiding prejudicing the investigation. This is because it casts the problem as a modal one. If we look again at the examples I provided, we can preserve the idea behind something being dominant or central, or alternatively, marginal, without going so far as to allow that these can only be understood in terms of *always* and *never*. And in the next chapter, we will look at a range of particularists who do exactly this, who try to say something general about what is centrally important without using the concepts of always or never.

1.10 The Correct Way to Understand the Flattening Objection

To sum up, the three theses I defend about the flattening objection are as follows:

1. The flattening objection is directly the central problem for ontological particularism, and indirectly, for a particularist viewpoint in general, with the latter inspiring the former.
2. Various other apparent criticisms of particularism can therefore be reduced to the flattening objection.
3. The discussion seems to imply that particularists must posit at least two "moral kinds" to allow for the flattening objection.

All of this points to a provisional conclusion that particularists might make their position more convincing by focusing on *existence*, and specifically, what is the ontological difference between what is or ought to be *common*, and what is or ought to be *rare.* Dancy, who has written the most on different forms of particularism, modified his outlook in view of some of the above responses to his original books and papers that introduced particularism into the literature. In a discussion with philosophers Johan Brännmark and Andreas Lind, Dancy discussed the evolution of his view, which, as we shall see in the next chapter, was motivated by considering the flattening objection:

> [Brännmark]: “One accusation that has been made against your views is that of vagueness. Would you say that part of the reason for that is that particularism is a spectrum of possibilities or is there something else?”
>
> [Dancy]: “Sometimes this accusation just means that I don't give any solid guidelines, which is true but hardly a complaint – it is just a description of my position. But sometimes people say that the view is essentially vague in places where it ought to be precise, and that is more worrying.” (Brännmark and Lind, 2008, 8)

This is indeed just a description of his position, which as I have noted, is what helps to set up the flattening objection. He later clarified what he understands the apparent vaguery to be:

> [Lind]: “In *Moral Reasons* (1992) particularism is a position in moral epistemology and in *Ethics without Principles* (2004b) you say it is a position in moral metaphysics.”
>
> [Dancy]: “I now take the view that there is a distinction between being a reason why one shouldn't do something and being a reason for believing that one ought not to do it. The latter topic is investigated in moral epistemology. “How do you know that it is wrong? What reason do you have to believe that it is wrong?” The former investigation is moral theory proper, where you investigate what makes an act right or wrong, which isn't an epistemic matter at all. If you ask which of those two are primary, it is clear to me that it is what it is for a feature to make an action wrong or what goes to make an action wrong, and the epistemic question is secondary. So, I think of the matter as what makes the wrongness, that's a metaphysical question, and how we track this is the epistemic question, the answer to which should largely be informed by the answer to the first question.” (Brännmark and Lind, 2008, 11, dates added by me)

The point here being that particularism’s necessarily most basic claim is, at first glance and according to the usual way of doing analytic metaethics, the one about *existence,* not about *knowledge.* What *makes* the action right or wrong is what the flattening objection ultimately wants us to consider. This is where their view is apparently deficient, and on realist philosophy, all of the more specific practical problems come downstream of this, and for this reason.

1.11 Globality and Locality

With everything we have seen so far in hand, there is a distinction I wish to draw in relation to the flattening objection, and what metaethical theories ought to be like. The moral landscape metaphor allows us to recast the distinction between the general and the particular as one between what is *global* and what is *local*, and this is the terminology I shall prefer going forward. On generalism, one takes the entire map of the moral landscape in one's hand: on the deontic theory of Ross, once again, once one knows the seven basic moral truths, one knows the entire moral landscape. One can go anywhere and understand one's surroundings.

On particularism, there is simply no concept of the global, or of having a complete map. There is only the local, where one can see one's immediate surroundings, and has to decide about where one is, or where one should go, according to this information only. The concept of a global map is simply required, the generalist seems to contend, to explain all sorts of moral phenomena; to explain education, or to explain the proper functioning of a society. The examples about social order and intrapersonal justification in particular seem to support this need for a global, and therefore shared, map. If the particularist takes this challenge seriously, there are two obvious directions they can go. Firstly, they could attempt to incorporate some understanding of globality into their position, while attempting to preserve its spirit.

Secondly, and more ambitiously, they could attempt to show how the concept of the local alone is sufficient to explain all kinds of moral phenomena. In either case, these explanations, which putatively ought to be metaphysical in nature, have to keep in mind that the practical purport of the flattening objection is properly an epistemic one. What we want to know in the end is how people decide what to do, or justify those decisions to others. If they choose to attempt to incorporate some kind of global understanding into their view, they will be adding to the metaphysical part to support the epistemic part. If they more ambitiously aim to only keep the concept of the local, they need to explain how the epistemic part actually does away with the substantial need for the metaphysical part. In other words, they can aim to solve the flattening objection, or dissolve the force of the objection.

1.12 Metaethical Correspondentism

There is one final term I wish to introduce before we move on, which is metaethical correspondentism. This term was coined by philosopher Kevin DeLapp to describe what he

identifies as the prevailing methodology in recent metaethics. On correspondentism, we don't start from nowhere, but instead we take ordinary moral discourse and behaviour, naïve views about metaethics, and "common sense" as our data, our starting point for metaethical theorising (DeLapp, 2023). We try to develop metaethical theories that explain, and *correspond* with what we take ordinary life to be like. I have deliberately left it until the end of this chapter to introduce this term, precisely because I think that *every* philosopher I have discussed takes such a correspondentist approach to metaethical theorising. They take their worries about anything from moral education to shoelace colour to be obvious expressions of common sense, and, on correspondentism, we think that metaethical philosophers are obliged to take note of such.

Anyway, in the next chapter, I will begin in the most modest way: by looking at all the ways particularists have been able to incorporate the concept of a global truth into their view. Presumably, this would be the easiest way to deal with the flattening objection. And there are nine such ways. Since particularism is characterised as an antitheoretical, negative view, adding extra *positive* pieces of *theory* to the view is held to *weaken* it. For this reason, all views that aim to solve the flattening objection by incorporating the concept of a global truth are called "weak particularist" views. I will explain why I think they don't work; why I don't think any of them, including Dancy's, provide a satisfactory solution to the objection as set up.

2. Responses to the Flattening Objection

This chapter considers a range of ways particularist philosophers have attempted to meet the moral non-flatness requirement, which again, was defined by Strahovnik as follows:

> The moral non-flatness requirement says that any moral theory must somehow account for the fact that some considerations or features of acts are more central to morality than others…particularism cannot properly account for the thought that some considerations are more central than others. (Strahovnik, 2016, 67-68)

The positions I consider are those that focus on the following concepts or approaches respectively: default reasons, semantic generics, defeasible generalisations, evaluative generalism, thick concepts, virtues, subjective contexts, and ideal theory. All of these are variations of "weak" particularist views, in that they are held to *weaken* the usual particularist starting position, which is the steadfast claim that there are no generalities in ethics, by attempting to include *something* that is general or constant, that is carried between cases, yet falls short of a substantial general principle. As I said before, all of these are attempts to *solve* the flattening objection, when that objection is cast in metaphysical terms, in that they incorporate some extra ontological content. This extra ontological content, in the vocabulary set up previously, may do either or both of the following: make a modal claim about the way things *often* are existentially, and/or make a claim about the way things are globally, rather than locally.

It may be worth briefly distinguishing weak particularism from weak generalism, which is also discussed in this literature, and is superficially similar, but really importantly different. Weak generalist views are those that seek to maintain the concept of a general law, rule, or principle, but allow for limited exceptions or excuses to be made. On these sorts of views, we can say that while there is still a moral law that lying is wrong, in suitably pressing circumstances, lying may be condoned on a case by case basis. While developing a view like this is indeed another way one can attempt to mediate between general and particular thinking, to incorporate the notion of context, they are not vulnerable to the flattening objection in the way particularist views are, and are ultimately meant to be part of a rejection of particularism anyway.

The challenge for weak *particularist* philosophers is to say something general about ethics *without* going so far as to invoke the notion of a law, rule, or principle. These two families of views should be taken to be approaching the middle from opposite directions, as it were, just as in politics we have the concepts of hard right, hard left, moderate right, and moderate left. What defines a weak particularist view is that it aims to get as close to the middle as possible while somehow preserving a rejection of moral laws. What defines a weak generalist view is that it aims to get as close to the middle as possible while somehow preserving the moral law. And the flattening objection is aimed at, and only weighs upon, one half of this spectrum, the motivation for the weak generalist being instead to incorporate some sympathy to context into an otherwise traditionally generalist view, and therefore apparently otherwise preserving the status the general principles are meant to have.

2.1 Default Reasons

The most popular way particularists have chosen to respond to the flattening objection by far is by invoking the notion of a default reason, and, as we shall see, several of the other sorts of responses seem to reduce to the default view in some sense. A default reason is one that, in the absence of confounding factors in some unusual context, has a default status of being "switched on" or "switched off". So, the fact that an action will involve telling a lie, for example, will *prima facie* count against performing that action, unless some other, stronger reason, or disabling feature, is present. We can readily see how this concept, if defended and accepted, would straightforwardly solve the flattening objection. A list of central concepts in ethics can be identified as default reasons, which have some kind of central or dominant status, while more marginal information can be given a "nondefault" status, where it has no particular valence outside of any particular case, and even then, perhaps none.

While the notion of a "default reason" in something like this form can be found in the logical and epistemological literatures from the 1970s onwards (e.g. Reiter, 1980) in general discussions relating to the formation of beliefs, it's very first appearance in any specifically *moral* literature was in 1999, where Dancy writes, here agnostically:

> The notion of a default reason - a feature that is a reason, as we might put it, other things being equal, or unless something prevents it from being the reason

> it would otherwise be. Now I myself do not know whether particularism should set its face against the possibility of such reasons. (Dancy, 1999, 27)

Later, in *Ethics Without Principles,* he gives a fuller definition of the concept:

> A default reason is a consideration which is reason-giving unless something prevents it from being so. The idea is that some features may be set up to be reasons, in advance as it were, although it is always possible for them on occasions to fail to give us the reasons they are set up to give. One can express this idea more or less metaphorically. More metaphorically, one could say that some considerations arrive switched on, though they may be switched off if the circumstances so conspire, while others arrive switched off but are switched on by appropriate contexts. Less metaphorically, one could say that if a default reason-giving feature does give us a reason in this context, there is nothing to explain; we only have something to explain when such a feature doesn't provide a reason. With other features it is the other way around; if they do provide reasons there is something to explain, and if they don't there isn't. (Dancy, 2004b, 112-113)

Therefore, Dancy himself eventually sought to weaken the original "strong" particularism he had introduced into the literature, in response to sustained criticism of the sort I have focused on in this thesis. What is interesting about this definition is that it is given in terms of there being nothing to explain when the reasons have the usual status, but that an explanation is required when they do not. This seems intuitive, which can be illustrated by considering the following two examples. It seems sensible that if a violent attack is wrong, no special explanation is required. But if a violent attack is held to be right, we might be curious why *this particular* violent attack ended up being morally right. Similarly, it might be no mystery if shoelace colour is irrelevant to some moral decision, but if it ends up being centrally relevant, we would be curious why.

Simon Kirchin provides a slight further gloss on this introduced concept and its definition in terms of *explanation*, by way of providing adding a further detail of how these reasons get to be this way: how they come to be switched on or switched off. A default reason

is one which requires only a single morally relevant fact to ground it (e.g., this act would involve telling the truth) whereas non default reasons require some further contextual grounds:

> Our main reason for denying that colours and days of the week have default valency was precisely the fact that by themselves they could not generate reasons. Even if pleasure, say, on this conception will seem as if it can switch its valency often, it might still have a default since we assume it to be able to generate reasons on its own. (Kirchin, 2007, 26)

This point about colours and days of the week not being able to generate reasons themselves is a reference to Dancy's idea about the difference between the grounds of reasons and secondary enabling factors, which I will explain here, since we will return to it throughout the thesis. Dancy's understanding of context rests on the idea that certain natural facts (e.g., that an action causes pain) are the ones that actually *ground* moral facts, whereas other "secondary" contextual background factors (e.g., that the action which caused pain happened during a wrestling match) affect our moral conclusions, without themselves directly being grounds.

Dancy divides these secondary factors into four kinds: enablers, disablers, attenuators, and intensifiers. Enablers and disablers, as the name imply, enable or disable some fact to be, or from being, morally relevant. As in the above example, presumably if someone feels pain during a consensual wrestling match in a fair sporting competition, this would "disable" this fact from making the match immoral. Attenuators and intensifiers weaken or strengthen the importance of some particular fact, respectively (Dancy, 2004b, 38-43).

Anyway, the metaphysical definition of a default we arrive at, then, is apparently a reason where an explanation is never required for it to have its usual valence, and always required for it to have the opposite valence. Default and nondefault reasons will apparently be the two metaphysical, moral kinds posited by the formulation of the flattening objection. But what this definition does not do is tell us why those two kinds have those natures. Instead of going "all the way down", this definition sits at a penultimate level. Presumably the ultimate "why" about the moral reality would let us understand what these things are in the first place. As Bakhurst, who wondered whether particularism can really be squared with the "true nature of morality", points out:

> Dancy invokes the notion of a default value in order to rebut the objection that particularism leaves nothing constant in the realm of value… Dancy thinks the notion of a default helps meet a challenge I raise in Bakhurst (2000), where I argue that the particularist needs to explain how holism is consistent with enduring moral commitments (e.g. against torture, for charity). In my view, however, the notion of a default is presently so under theorized as to be useless. Dancy appeals to an explanatory asymmetry...What we need to know, however, is what, in the case of default reasons/values, grounds these explanatory asymmetries." (Bakhurst, 2005, 278)

Our real question, then, is "what makes a default a default?" And the first way philosophers have tried to ground default reasons is in descriptive features of our world. In realist moral philosophy, the moral facts are consistently "grounded" in the natural, empirical fact, and the idea here being that the way our world is *naturally* guarantees that it throws up overwhelmingly and repeatedly consistent sets of grounds, and hence, evaluations. This might be what makes those "central" facts *often* so important in a way that is nonarbitrary and nonaccidental.

While Dancy never ended up saying anything more basic about defaults, and confined himself to simply introducing the possibility of the same, Pekka Väyrynen (2004) and Kirchin (2007) both identify an obvious and specific way of descriptively identifying these sorts of central facts in our world: statistical generalisations. This approach involves surveying a great range of actions, situations, contexts and drawing some conclusions about what seems to be so *regularly* and *overwhelmingly* the case. When applying this method, for example, we might look at a great range of actual cases where somebody said something that was not true, and conclude that in the majority of the cases, it was the wrong thing to do. But Väyrynen finds this sort of approach ultimately unsatisfying, concluding:

> If default reasons are to be at all relevant to the generalism-particularism debate, they must be construed as lacking precise implications concerning their extensional generality, by which I mean the ratio of cases in which reasons are undermined relative to those in which they aren't … their extensional generality depends on the sorts of contingent worldly matters which cannot ground moral reasons. Default reasons cannot therefore be recruited to provide any further support for particularism. (Väyrynen, 2004, 54-55)

His argument for this, based around a thought experiment about a possible "nasty world", is a modal one:

> A Nasty World is a world where – whether by nature, nurture, or both – human interaction is coercive and fraudulent and human psychology sadistic, deceitful, and perverse. The endeavors of the denizens of such a world would provide no rational basis for a general expectation that promises, enjoyments, and other considerations are normal with respect to reason-giving, since they would normally be undermined, and so would normally not give reasons. Manipulative situations, for example, would be in the relevant sense normal with respect to promise-making – these are the sorts of situations in which promisors in a Nasty World tend to find themselves, thanks to the general shape of their world. Since manipulated promises, at least qua promises, are abnormal with respect to reason-giving, it follows that in a Nasty World the fact I promised to ϕ wouldn't normally give me a reason. (Väyrynen, 2004, 63-64)

The point here being that, on the statistical generalisations view, the only way we can understand all of the popular modal words and phrases used in this debate (often, more often than not, usually, standardly, most of the time…) is statistically, which means that something ought to have obtained at least 51% of the time to be considered a "default" outcome.

It is a matter of definition that "most of the time" or "more often the not" must be construed in this way when understood statistically. If the same thing had instead obtained 49% of the time, we would have to conclude in favour of an opposite norm, to the point of absurdity. For example, if telling the truth is right 51% of the time, then the default is "lying is usually wrong". If that number drops to 49%, the default is "lying is usually right".

But charitably, *prima facie* we don't think that the central norms we *actually* have in *this* world are on such thin ice; we are not so suspicious or sceptical about the way things actually are. The whole reason we even have a flattening objection is that we do seem to think that lying is *usually* wrong in our world. So instead, I would prefer to develop Väyrynen's point about "Nasty World" by instead continuing to focus on our own world, since I think the point can perfectly well be made there.

And this will involve considering precisely what we mean by a "context". Particularists claim there are an open ended amount of natural facts, arrangements of, and evaluations of, the

same. One reason for this is that the same familiar natural facts are endlessly evaluated anew each time we come across them (e.g., that something is painful) depending on the context we find them in, which other facts are present, and which secondary factors are present. A second, perhaps more interesting reason is that new natural facts are constantly being created that often have great moral relevance. The first sort will often end up being evaluated differently because of context, and the second sort, more strongly, are part of completely new contexts actually being created.

I take the first sort first. On the statistical generalisations view, whether we consider Väyrynen's theoretical "nasty world", or simply our own world, what we might call ordinary life, contrary to what may perhaps seem intuitive, actually can be enough to problematise the trend. We don't need extraordinary events for telling untruths, for example, to become just common enough to become unevaluable by terms like "often" or "overwhelmingly so". But more strongly, extraordinary events that we know in fact regularly take place will create contexts where certain default reasons seem to so often not apply. During wartime, for example, as in World War One or World War Two, the fact that something may involve violence, or may involve saying something that is not true, can regularly be facts about actions that we nevertheless regularly end up thinking we ought to perform nevertheless.

A straightforward example is provided by the fact that during World War Two, some European people took it upon themselves to help hide Jews from the Nazis, or to help them to escape. A similar example comes from the century before, in America: some Americans took it upon themselves to help slaves escape to freedom, as a reaction to the context, which was the institution of slavery, which was normal at the time. So descriptive considerations of what is either known to be ordinary or extraordinary in our world as we have known it so far are already enough to wonder whether the default reason view really has the basic, ultimate explanatory power it might need to have.

Secondly, another important definition of "context" to consider is that of a completely new or emerging future context. The first sort were where currently existing natural facts are arranged differently, or evaluated differently. This second kind is where new *natural* facts are introduced into the world that literally didn't exist before. One example is manmade climate change. Further relevant examples are the invention of nuclear power, or more importantly, the invention of nuclear weapons. The natural facts are the empirical ones, those that are studied by the natural sciences. And again, in realist philosophy, the general moral truths are always grounded in the same natural facts. The generalist realist viewpoint needs these facts to always

be the relevant ones, and to always be available. But this is obviously not true even at the scientific, preevaluative level: the level where the same familiar set of natural facts is there to be evaluated once.

If, for example, it came to pass that the Amazon rainforest had been completely cut down, or that an exchange of nuclear weapons had begun, or that melting polar ice had caused sea levels to rise severely, what new moral contexts and, and therefore decisions to be taken, would be created? And how much confidence can we put in our beliefs about "the way things usually go on earth"? Examples of this sort can be given *ad nauseum*. It is currently popular, both within and without scholarly circles, to make proclamations about how the development of artificial intelligence will radically alter human life. Some scientists claim that humans will eventually colonise Mars. Others claim that human lifespans could be heavily extended, or even that immortality could theoretically be achieved. These points about the generation of completely new natural facts that need to be evaluated is my own contribution to the particularist debate, since particularists so far have usually focused on the first sort of understanding of context: the particular arrangements of already existing natural facts, and I think the addition of this second sort of definition only strengthens the particularist thesis.

Anyway, by taking these two definitions together, we seem to end up where we started. Thinking carefully about the way our own world actually is and has been doesn't seem to provide some obvious central truths, and accepting our world will constantly change means accepting the future may not be like the past. With this conclusion in mind, the particularist viewpoint seems to rule out at least the default reasons view being a substantial solution to the flattening objection. I think Väyrynen is right to call such reasons "contingent", and in the terminology I set up earlier in the thesis, such generalisations will only, at most, hold "locally", and will say nothing guaranteed about what is often the case globally, which is what the flattening objection seems to require. Therefore, some other solution will be required. So next, I consider a similar sort of weak particularist view that is at least meant to avoid some of the problems the default one has: one based around the concept of "generics".

2.2 Generics

Another strand that runs both within and without the moral literature is one about "generics", sometimes "semantic generics". Like the concept of a default reason, the concept of a generic

statement can be found in a range of literatures from the 1970s onwards, and predates the particularism/generalism debate (e.g. Lyons, 1977). A generic is any general statement of fact, about any domain, which is accepted as true in a discussion in that domain, by the relevant speakers (e.g., experts, participants) in that domain, according to the standards of importance, truth, or relevance in that domain. By a domain we mean some specific area of inquiry, or topic of discussion. On this definition, moral philosophy, physics, wine tasting, and dog breeding are all "domains". In the context of the flattening objection, the generics approach involves defining the central ethical truths in terms of their being accepted as valid generic statements about morality.

Crucially for our purposes, considering my conclusion that the default reasons view was insufficient to solve the flattening objection, the generic view is meant to preserve the spirit of the default definition we started with, in a way that the statistical generalisations explanation of the same cannot. Ravi Thakral has done the most to try and develop the generics approach specifically in terms of explicating moral particularism, but references to it have been made by a range of particularists (e.g. Lerner and Leslie, 2013, Lance and Little 2004, Little, 2001b, Väyrynen 2004).

The generics view involves pointing out a variety of statements from different domains that are held to be both true and explanatory, while in fact only holding sometimes. The statements are modal, but their meaning or truth does not rely on this modality being explainable statistically, in terms of actual prevalence. And of course, this need for statistical explanation in terms of actual prevalence was what we found problematic about the defaults view. This means that Thakral's approach directly meets the challenge Väyrynen gave to those who otherwise find defaults promising, which was, again:

> If default reasons are to be at all relevant to the generalism-particularism debate, they must be construed as lacking precise implications concerning their extensional generality, by which I mean the ratio of cases in which reasons are undermined relative to those in which they aren't. (Väyrynen, 2004, 54)

The understanding of defaults as requiring actual prevalence was what I identified as specifically problematic about the view, and Thakral's first explanation of why generics are unalike to default generalisations in this sense is as follows:

> Whether or not a generic is true need not correspond with statistical regularities. 'Ducks lay eggs' is widely accepted as true; yet, it is only female ducks of reproductive age that lay eggs. Thus, a majority of ducks do not need to lay eggs in order for the generic to be true. It is widely accepted as true that 'Ticks carry Lyme disease'; yet, only around 1% of ticks are infected. That is, we accept the generic even at such a low prevalence level. At the same time, there are generics that are unacceptable even when the majority of a kind possess a property. We do not accept 'Books are paperbacks' as true despite the fact that the vast majority of books are in paperback form. (Thakral, 2023, 208-209)

Thakral provides a further list of examples of generic statements:

> a. Koalas sleep all day.
> b. Americans drink coffee.
> c. Tigers are striped.
> d. Generation Z cares about the environment. (Thakral, 2023, 206)

All of these statements are held to be meaningful despite being statistically undetermined. Thakral goes on to provide a pair of definitions apparently drawn from the usual meaning of the examples he considers:

> 1. Resistance: Generics can remain true even if there are counterinstances.
> 2. Non-Numerity: Generics do not carry any information about how many instances are required in order for the generalization to be considered true. (Thakral, 2023, 208)

I list these because by observing the progression we can clearly see the two relevantly central notions here. The supporting statements presented are firstly all of the "companions in guilt" sort, where a certain form of understanding can be taken from another domain and brought over to be applied in ethics, and secondly, are apparently accepted as true in their respective domains regardless of how general they really are, or alternatively, as in the books example, rejected, regardless of how generally true they may really be.

Since this is the first time I mention the concept of a "companions-in-guilt" argument, and I will return to it regularly throughout the thesis, I will briefly explain it here. Companions in guilt arguments are arguments where, firstly, some relevant similarly is drawn between questions, claims, arguments, or conclusions in two different domains, and secondly, it is claimed that if we accept one, we must accept the other. This is meant to be because of the relevant similarity between both the claims and the domains. The comparison can be made between philosophical and nonphilosophical domains, or between two philosophical domains.

In this particular case, the idea will be that if we accept generic statements about every other topic as being true and meaningful, why wouldn't we accept this for generic statements about morality? Presumably if we would reject the status of such statements *only* for morality, there must be something importantly different about morality compared to any other topic of discussion. Otherwise, we would be inconsistent. The "guilt" part of the companions-in-guilt approach refers to rejection instead of acceptance, and it holding in the other direction. In other words, if we dismiss generics here, and *not* for reasons that are specifically only about the nature of morality, we might ought to dismiss them everywhere (Lillehammer, 2019). Thakral's approach, then, in common with any companions-in-guilt approach, relies on our independently accepting that moral speech is relevantly similar to nonmoral speech.

Anyway, the first of Thakral's definitions, resistance, is not what is unique about his view. The default reasons view, and in fact, all of the sorts of weak particularism I will look at, of course allow for counterinstances, and therefore incorporate "resistance". The second definition, the one in terms of non-numerity, is both what is interesting here and the one that ought to give it an advantage over the default reasons view, where a key problem was that a default was meant to explain what is "often" or "usually" the case, and we weren't sure whether it really does. So Thakral is attempting to explain how general ethical statements can still be meaningful while jettisoning the baggage of also having to commit to their obtaining so very regularly.

I will take both of these features together in my response, the companions-in-guilt methodology and the conclusion of non-numerity, since I think pointing out what is wrong with a "companions in guilt" type approach here also helps us see what might be problematic about accepting non-numerity here. And there are at least two things we can say about these other domains that form the companions-in-guilt. The first problem I have with generics is that of seriousness, and the second the one of kind.

I consider seriousness in two ways. Firstly, internally to the moral domain. At least one uncontroversial way we can parse what is meant to be "central" vs "marginal" in ethics is by perceived seriousness. It would be a strange response to the flattening objection that didn't suggest that "torture" or "murder" are central deontic concepts at least partly *because* these are so very serious topics, or that justice or cruelty are central evaluative ones. This is very much predicated on the intuitive, apparently practically grounded nature of the debate. Again, both particularist and generalist metaethical philosophers are usually metaethical correspondentists. They think our metaethical views ought to fit with our experience of ordinary moral life. On correspondentism, any serious metaethical view will allow that murder and torture are centrally important moral concepts.

Secondly, and more importantly, how seriously we take *different* domains. A key part of morality is how much we care about it, and care about getting it right. Apropos the point made just above, about centrality, the more central we take a question to be, the more we care about accuracy still. While we are perfectly happy with "tigers have stripes", knowing that some don't, presumably we wouldn't be as happy with "all thieves should be executed", while believing that some shouldn't. So, the point is that we both take morality to be more serious than other topics, and that we take some particular bits of morality more seriously than others.

The discussion of "rule fetishism" in relation to particularism by Andrew Gleeson helps to highlight the problem of seriousness. Gleeson cautions us that strict adherence to rules can regularly lead us to severe moral failure, if we insist on always applying them no matter the context (Gleeson, 2007). To repeat an example I mentioned earlier: if it is World War Two, and my neighbour helps a Jewish person escape from the Nazis, my reporting him because of my insistence on cleaving to the principle "always tell the truth" has not only lead me to moral failure, but lead me to moral failure precisely *because* it was done in the name of respecting moral truth. The mistake is made twice over.

The point being that principles don't just lead us to make bad choices (which we can also do perfectly well under particularism) but more strongly, that they actively compound and exacerbate our chances of making a bad choice. In the example above, the person who tells the truth to the Nazi's has actually actively been *hindered* by possessing a moral principle; it would have been better if he didn't have one. So how does the semantic generics view look in light of rule fetishism?

None of the examples Thakral gives would cause much harm were they believed completely strictly or obsessively. If there is a man who, contrary to exceptions presented to

him, believes fanatically that all tigers have stripes, or that all ducks lay eggs, we won't be overly interested in criticising him. But the same is simply not true for the moral domain. Rule fetishism is held to be able to do real harm there, and I think this is part of what motivates such intense philosophical interest in realism, and metaphysical models: we really care about what the answer is! [2] Consider the example of a man who obsessively and repeatedly claims that "most black people are criminals". We wouldn't take this so lightly: we would want to strongly criticise him, because we know how much it *matters* whether such claims are true or false. And this is precisely because we take the weight of this claim *seriously*. We know it can cause genuine moral harm. Morality is a realm where we deeply care about the actual accuracy of our moral beliefs and moral claims, not generally or for the most part, but actually in each and every case.

So, companions in guilt arguments rely on *relevant similarities* between the domains compared, and I think my comments about seriousness highlight at least one possibly relevant dissimilarity. But there is another way the various domains discussed differ, which is in subject matter, and this brings me to my second point. All of the other domains that feature in the generic examples provided by Thakral are natural ones, that are identified empirically. Natural domains are familiar, completely understood (for our purposes) and uncontroversial. A popular example of a generic given in the literature is "all birds can fly", even though penguins in fact cannot fly (e.g. Reiter, 1980). The kinds, populations, and flying abilities of birds are in fact no mystery to us and are about an uncontroversial *kind*. Ornithologists have comprehensively studied all the kinds of birds, and there is no dispute about whether penguins can fly. It has been proven and agreed they cannot, in the way that is relevant to the natural sciences. But the "nonnatural" moral domain, which the particularism/generalist debate is about (both being versions of nonnaturalism) is precisely not usually understood like this.

Mackie developed the influential "queerness" argument against the concept of a nonnatural moral kind, claiming that such an object has a "queer" (strange) character (Mackie, 1977). Many philosophers have separately questioned how we can have knowledge of such kinds. There is no agreement about how many, or any, moral kinds exist, or how we can or do even know. In the debate about moral particularism and potential flatness, the mysterious nature of moral kinds and our apparent lack of knowledge of the structure of the moral domain

[2] Realist moral philosopher David Enoch (2011) actually provides an argument from seriousness specifically *for* moral realism, arguing that taking morality seriously requires us to find it to be real.

is precisely what is held up for discussion. These are just some of the reasons why I think it won't be straightforward to justify moral generics by finding companions in other domains.

In conclusion, I think because of the unique character of moral thought, and the seriousness of the same, a companions-in-guilt type story about semantic generics is of questionable utility. Further, I think that numerity cannot be so easily discarded here as it might be when discussing tigers, or coffee. While it shouldn't be a received truth that *only* a statistical definition can give weight to central moral concerns, we shouldn't go too far in the other direction either. The flattening objection is meant to be about what is really central to moral life, and it would be surprising if the solution was one that meant what is central might also almost never be true, or relevant. Väyrynen suggests, and Thakral concurs, that we should be relaxed about how often "central" moral truths do actually obtain. I am sceptical about whether this is the right direction to go in. Here, I am inclined to echo Dancy, who wrote, on the topic of companions-in-guilt arguments:

> I just don't think that much is gained by finding a view similar to one's own in a cognate area which has had a good run for its money there. This might achieve respectability for one's own view (were this in danger of being denied) but could hardly count as evidence for its truth. (Dancy, 1992, 94)

This "respectability", I think, is generally what all companions-in-guilt type arguments hope to achieve, and what Thakral is looking to do here. As a rhetorical strategy, pointing out how we commonly use language elsewhere can at least seem to make a similar analysis of apparently similar statements in ethics seem more respectable. But ethics has a specially important character, which leads me to think that when responding to the flattening objection here, we ought to provide some solution which stands up based on ethical considerations or arguments alone.

2.3 Defeasible Generalisations

A third approach to solving the flattening objection, which has something in common with both the default reasons view and the semantic generics view, is one predicated on the concept of a "defeasible generalisation". Such generalisations were introduced into the particularist debate

by Lance and Little, and like particularism itself, were originally inspired by considerations about epistemology, about correct belief forming practices, and later extended as a view about specifically moral beliefs. Lance and Little describe their epistemological position, and the analogous moral view they derive from it, as follows:

> There are many generalizations in epistemology which, though full of exceptions, seem essential to the very enterprise of knowledge. "The future will be like the past," "things are as they appear to be;" "people tell the truth… In none of these cases is the generalization universal or exceptionless… In this regard, contextualism in epistemology seems analogous to claims made in ethics by a view known as "moral particularism" … But comparisons to moral particularism should give us pause. If claims of contextualism are widely regarded as uncontroversial in epistemology, the claims of moral particularism are widely regarded as controversial at best, downright crazy at worst… If these problems are serious in ethics, they should be for epistemology as well… For both disciplines, strategies for answering these queries divide into two broad camps. "Radicals" embrace – indeed, celebrate – the claimed absence of law-like structures…the justification of a given belief, it is said, is a matter of discerning how the moral or epistemic considerations add up in each individual case… "Moderates," in contrast, argue that this misreads the lesson of contextualism. The presence of exceptions does not mean we have left theoretical generalizations behind, but that… we must soften their semantic content and recognize their nature as "ceteris paribus," "hedged," or "defeasible" generalizations: defeasibly, pain is bad making; defeasibly, appearances are to be trusted. We believe the second camp has to be the right one… We argue that there are defeasible generalizations whose semantics must be understood in terms of a normative conception of "privileged conditions". (Lance and Little, 2004, 435-437)

Once again, we find a view alike to moral particularism, but here about epistemology, described as "radical". And this radicality, of course, is what the responses to the flattening objection we survey here, including Lance and Little's, are meant to ameliorate. Anyway, for Lance and Little, these sorts of generalisations are grounded in privileged *conditions*. While some

participants in the debate (e.g. Väyrynen, 2023) use the terms “default” and “defeasible” interchangeably, and seem to take the two positions to just be the same thing, I think this predication on conditions is what distinguishes Lance and Little’s version, since the default reasons view was predicated on statistical trends, and the semantic generics view was apparently predicated on a proposition being accepted as relevant by participants in a particular domain.

Indeed, they begin by referring to the default reasons view and the semantic generics view as the two currently available in their explanation of how their own view is different from both, and improves upon both (Lance and Little, 2004, 438). Like the semantic generics view (and indeed, others) the defeasible generalisations approach involves a companions-in-guilt style comparison between two domains: although the comparison here is instead between two domains in formal or theoretical philosophy (epistemology and metaethics), rather than between philosophical domains and entirely nonphilosophical ones. I will not repeat my thoughts about such comparisons here, but instead attempt to focus entirely on the concept of privileged conditions in moral contexts, since perhaps that can stand up on its own.

Lance and Little identify two characteristics of a privileged condition. The first is classificatory dependence, and the second is justificatory dependence. They develop both with a range of examples, mostly from nonmoral domains. Classificatory dependence is defined in terms of a concept forming part of a class, where there is an acknowledged “main” member of the class, and the other members of the class are understood to be particular derivations of the main member. This main member is the one that is “privileged”. Using the example of football, they point out that “normal” football is played with 11 players, with a particular goal size and pitch size. Other forms of football, like five-a-side football, are understood as football precisely because they both differ from the main example but endorse the spirit of its features.

This is what is meant by the main example being part of the explanation of why something different should still be understood in the same way. It only makes sense to say “five-a-side football” precisely because eleven-a-side football exists (Lance and Little, 2004, Strahovnik 2016). On classificatory dependence, the main concept is not just what the others differ from, but *always* part of the explanation of why they so differ. Further, the contextual difference is nonarbitrary and readily understandable. The pitch and goal are smaller in five-a-side football, because the team is smaller. Lance and Little call these “riffs” on the central “theme”.

Lance and Little also use the example of ornamental chairs, which are not actually for sitting on. These are part of, and depend on, the class "chair", and are understood as chairs even though they cannot be sat on, precisely because chairs usually can be sat on. Further, ironic or sarcastic speech is meaningful precisely because we have the "normal" concept of serious speech, which they are variations of, and are compared to. In all cases, the main member of the class sets certain "conditions" which members of that class exist within, which, as I pointed out, is how Lance and Little distinguish their view.

Justificatory dependence is explained according to a justifying/sceptifying dyad. Lance and Little use the example of a red cup. In "normal" conditions, one is justified in believing that because one sees a cup as red, it is red. One should only be sceptical if the *conditions* are in some way modified, e.g. the light level is low, or one has taken strong drink, or taken drugs. Again, this consistent focus on conditions is what is held to distinguish a defeasible generalisation explained in terms of dependence on such conditions, from a default, which is defined in terms of explanation:

> This point helps to illustrate an important difference between something carrying a defeasible import and operating as an epistemic default. To call something a default is to say that it is a justified "start here" position – an assumption one is entitled to make in the absence of special evidence to the contrary. Now we believe that defaults are essential to epistemology; in particular, we believe that they are essential if we are to be protected from skeptical collapse. For all that, though, they are thoroughly distinct from defeasible connections. A defeasible import is the import something has in suitably privileged conditions, the understanding of which is thought crucial to understanding the theoretical significance of the kind. (Lance and Little, 2004, 450)

Here again, it is made completely clear that the defeasible view is about *conditions* in particular. So, while I acknowledge that Lance and Little's defeasible view can be distinguished from a default one, or a generic one, and is therefore probably worth considering separately, I think it is still problematic.

To see why, we can refer to Strahovnik, who as we saw, also formulated the non-flatness requirement, and has picked up on and extended Lance and Little's concept of defeasibility. Strahovnik finds much to otherwise recommend the concept of defeasibility as

one that ought to be put to use across normative philosophy (he says that scholars in fields such as law and aesthetics ought to like the concept) but finds that it fails to meet his non-flatness requirement for particularism. And the reason for this is that he thinks it fails to provide any basic explanation; one that goes "all the way down". His own conclusion here is that Lance and Little's view fails, in the end, to be a form of particularism in that the riff/theme model both has and requires general principles as it's basis (Strahovnik, 2016, 64, 70).

So, does Lance and Little's picture built around privileged conditions, and dependence on the same, amount to an acceptable solution to the flattening objection? *Vis a vis* Strahovnik, I do think the key question is whether their picture is explanatory in the basic sense; whether it goes "all the way down". Just as with default reasons, where we asked "what makes a default a default?" we want to ask what really makes a member of a class privileged *in ethics*, or what makes a seeming specially justifiable *in ethics*.

On the first point, while apparently true, I don't actually think the notion of classificatory dependence is special to the defeasible generalisations view: I think the notion of a class is actually a more basic and general one in ethics, that the defeasible view is parasitic on, and needs to be argued for separately. This is because it relies on our knowing what the classes and privileged members are, and our agreeing where we get them from. One obvious answer is that the main classes will be "evaluative" and "deontic", with "good" and "right" being the privileged members, and all the other members being the riffs on the theme, as Lance and Little describe it. For example, surely the more specific evaluative concepts, which rely on the central one and are explicable in terms of the central one, will be kindness, justice, etc.

But accepting this just means returning to the beginning of what metaethics is, and to familiar ethical concepts, where we can say that one is more basic than another. This is just a definitional point. But what we really want to decide in our discussion of the flattening objection is how, on a case by case basis, we will decide what is kind here, whether truth is required here, etc. I don't think saying that our verdict that something is kind here is related to goodness really helps to explain it here. Neither does saying that our decision to be honest or not will be a riff on the concept of honesty, moderated by conditions, gives a basic answer here.

The epistemic version of the view (justificatory dependence) I will return to later in the thesis where I discuss epistemology, since I think it doesn't weigh on the ontological nature of the flattening objection like the classificatory view does. The classificatory view is about parsing different moral *kinds*. Recall my distinction (also made by Dancy, and by Bakhurst) between the structure of the external moral reality, and our skill in navigating that reality, with

the second being predicated on the first. The notion of justifying will be relevant to how we make decisions, but not what we make the decisions in regard to, apart from indirectly. To take the red cup example Lance and Little use to explain justificatory dependence: the relevant question *vis a vis* the flattening objection would be "why is the cup red?" or "what *makes* cups red?". Anyway, a useful verdict we can draw from Strahovnik's own criticism of Lance and Little (which I think is broadly correct) is that while it seeks to have a particularist character, the defeasible generalisations view is probably better understood as a form of weak generalism, as it really works on the model of basic rules and exceptions (Strahovnik, 2016, 64, 70).

There is a final general point I wish to make here, which probably applies to most of the weak particularist views I will survey, though I don't wish to keep repeating it, and suits the notion of privileged conditions well enough. And that is that particularism doesn't claim that duties *will* change between situations, just that they *can*, depending on the form of that situation, and this talk about conditions can also be taken to be the notion of a very particular situation repeated, in disguise. I will explicate this with reference to the duty of keeping promises. Imagine a happy, loving and trusting monogamous couple, where both of them have, and keep, high standards for how they treat each other, according to the standards that are internal to such loving relationships.

The wife promises the husband that she will not cheat on him. The first thing we can say is that, whether we are generalists or particularists, we will get the same answer: she should keep this promise because it is made against the background of the loving relationship, according to that background. But secondly, and more strongly, such a promise is superfluous and would actually not be expected, and this is because it is just taken to be built into, implicit in, the context of the happy relationship. This is another way we might understand the concept of "privileged conditions", or more broadly, any view about there being a general normative backdrop we make promises against, and keep duties in regard to.

But if we imagine that the relationship is instead a violent and hateful one, where the husband regularly lies to, neglects, and otherwise mistreats the wife, we no longer automatically have the same strong intuitions about there being background conditions she might make promises, or keep duties, in regard to or against. The context is quite different. And particularism is simply and precisely the view that duties only belong to specific contexts, with no expectation of how often or how regularly the same context will obtain. So we should be sceptical of any response to the flattening objection in the form of a weak particularist view that seems to refer to any kind of background conditions, simply because exactly the same

conditions and conclusions can be drawn on "strong" particularism. Lance and Little's understanding of a "condition" seems to have a relevant similarity with what we have been calling a "context"; and that is what we wish to explain.

2.4 Evaluative Generalism

Another way we might attempt to approach our topic is by turning our focus from actions to values: from deontic questions to axiological ones. After all, it is a truism in metaethics that any comprehensive metaethical or ethical theory will have firstly, something to say about deontic concepts, secondly, something to say about evaluative concepts, and thirdly, an explanation of the connections between the two, specifically in how our values inform our behaviour. At the level of ethical theory, the two main views, consequentialism and deontology, can be defined precisely by their focus on one or the other families of concepts, the importance of one or the other, and how each is related to the other, or derived from the other. On consequentialism, action is entirely informed by what is valuable, whereas on deontology, actions can be obligatory or forbidden without reference to value.

Particularism is a metaethical view which has no particular stance on such questions at the level of ethical theory, but perhaps saying something more concrete here, while remaining at the metaethical level, might be helpful in our discussion of the flattening objection. It may further be relevant to point out particularist writing is overwhelmingly deontic, and action centred. The key concepts that are focused on are reasons for action, the rightness or wrongness of actions, and practical decision making. While Dancy (and some others) have also outlined a particularist view of valuing (e.g. Dancy, 2003), particularism proper usually refers to the view that focuses on actions, and reasons for the same.

So, while accepting that particularisms insights are appropriate for actions, could we solve the flattening objection by separately accepting a more general understanding of values, in a complementary way? Might it not be the case that it is our values that are actually stable, and are the things we carry about from case to case? To borrow the terminology I set up in the previous chapter, our values will be more global, while our reasons for specific actions will be more local. In fact, we can even see a way where the addition of a general value actually seems to improve our particularist position about particular actions; where we must act differently here precisely *because* of some general value we wish to honour.

To my knowledge, only one person in the literature has mooted this view in the most general and bare form, e.g. that we should understand the axiological domain in the usual general way, and the deontic domain in the usual particular way, and aimed to predicate something about the second on the first. This is the view of Bakhurst, who, in the context of how particularists should respond to their critics concerns about moral education, writes:

> What I want is that my children should acquire certain values: there are some things I want them to care about when they are adults. I want them to be concerned about the wellbeing of other people, to respond to the sufferings of others and to take pleasure in their joy. I want them to respect others, and not to harm or exploit them. I want them to be just and to resent injustices. I want them to value equality, impartiality, and democracy. I want them to value special attachments to family and friends, but to extend their sphere of concern far more widely. I want them to be alive to the good things in life, to take pleasure in the pleasurable and to shun evil. I want them to be autonomous and to value autonomy. I want them to be prudent. I want them to be reasonable and to value reasonableness. (Bakhurst, 2005, 270)

While it is one that suggests itself as a possible philosophical position out of all the possible models and options, such a view would qualify as weak generalist, rather than weak particularist, and therefore wouldn't strictly be relevant for our purposes here. However, there are two more specific forms of this view that *have* been developed very much in the literature, and do qualify as weak particularist, and aim to respond to the flattening objection. And these are positions about thick concepts, and positions about virtues, about character traits.

2.5 Thick Concepts

The next way of delineating what is central that I will consider is by focusing on thick ethical concepts, whether evaluative or deontic. Thick evaluative concepts are more specific forms of goodness or badness, such as kindness. Thick deontic concepts are more specific forms of rightness or wrongness, such as murder. Thick concepts are an obvious choice of focus when considering some limited kind of moral information that might be both central and general. In

his paper about default reasons, Kirchin also summarised the various ways particularists have characterised thick concepts in the following taxonomy, apparently ordered increasingly according to strength:

1. In order for ethics to have a rational structure it is not necessary that there are any thick features that have invariable valence.
2. Not all thick features have invariable valence, some have variable valence.
3. No thick feature has invariable valence. (Kirchin, 2007, 56-57)

Now, the first of these is a claim about reasoning, or knowledge. This is a topic I will return to later in the thesis, but for now, I will again stick to considering responses to the flattening objection that are distinctly ontological in nature, which the second and third of this list are. The third apparently seems to be a restatement of what we expect the traditional particularist position is, while the second is the sort of milder view we want to investigate in this chapter: one which posits two moral kinds. I would like to introduce a 4th option to this list, which I think most accurately characterises what the opponents of "strong" particularisms ontological claim about thick concepts think it actually is:

4. Thick features *cannot* have invariable valence.

And *cannot* is the strongest of the possible options. It makes the second a nonstarter, while leaving the third insufficient insofar as it can be read as the sum of the rejection of various actual thick concepts, leaving open the possibility that a thick concept actually could have invariable valence.

Before we turn to consider which, if any, of these characterisations describe particularism most accurately, it is worth briefly considering again what invariant valence ought to mean. According to the usual schema of natural grounds and grounded evaluations, a thick concept would be a certain set of natural grounds and a specific evaluation that refers to those grounds. While each set of grounds is formally paired with the *thin* concept of value, the specific name it has refers to its natural qualities. An invariant thick concept would refer to that specific evaluation always belonging to that set of grounds, wherever they might be found.

Kirchin, in another text, makes this suggestion in a discussion about thickness and flatness, but does not do more than moot the possibility: "We are looking at things wrongly:

we might wonder whether mopping the torturer's brow, for example, is kind *in the first place*" (Kirchin, 2003, 68). This thought is in response to the question of firstly, on any view, if we say that mopping someone's brow is kind, and kind things ought to be done, would this apply to a torturer? Secondly, whether a particularist would want to say that it is simply not kind, or that it is kind, but also ought not to be done nevertheless. And this first angle is what I wish to more confidently develop. It is important to distinguish between specific and/or thick concepts, the actions they describe, and the particular moral evaluation of those actions. This is because I think particularists, including Dancy, actually under develop their own view in a way that seems to somewhat muddy the waters.

I think the correct understanding of particularism's understanding of thick concepts is *not* any of the 4 listed above, and has therefore been misunderstood by its opponents in a way that makes it seem far more controversial than it actually is. I will argue that the correct view of thick concepts that can actually be derived from particularism is:

5. All thick concepts have invariable valence, exactly as pretty much all philosophers understand them. But this is irrelevant for the purposes of solving the flattening objection.

This is because "that it will be helpful" or "that it would involve murder" are *not* the reasons for or against doing the action. Those concepts instead evaluate the action that will be performed. This is because ethical grounding involves the reason being *grounded* in the natural features of the situation at hand, and it is not the features that are valenced. The actions that are being performed are actually *assisting* or *killing*; those are the things that the particularist thinks are sometimes right, and sometimes wrong.

Helpful, and murder, on the other hand, are words that describe the relation between this rightness or wrongness and those facts, insofar as they are words we developed to specifically describe such pairings, not such actions. The definition of helpful is "assisting someone, and it is good to do so", while the definition of murder is "purposefully killing someone, when it was not right to do so".

Because killing is rarely justified, and helping others usually ends up being good, we turn these into the heuristics (or even rules) *be helpful* and *don't murder*, even though murder is a conclusionary word: "don't murder" is unusual because murder means "a killing you shouldn't have done, because it was wrong to do, or should not do, because you have identified

it will be wrong to do". [3] Ultimately, such concepts reduce to the general but unsurprising claim "don't do things that are wrong". "Not to be done" is built into any such thick deontic concept with a negative valence, so adding a further "do not do x" is redundant. What particularists would be interested in challenging is whether *killing* is always wrong, and whether there is such a moral principle as "don't kill!".

A further way we can illustrate this point is by pointing out that, at least on immoral actions that are sufficiently serious, we apply labels to the character of those who carry them out. For example, somebody who commits a murder is called a murderer, somebody who steals is called a thief, and somebody who tells lies is called a liar. But let us take again the example of somebody who, to protect Jewish people in World War Two, tells a lie to a Nazi. Do we call such a person a liar? At least in the usual sense where we wish to negatively evaluative someone's character, presumably we do not. And this is exactly the kind of situation particularism wishes to particularly make sense of.

The point here is that Dancy's position that negatively valenced thick concepts can count towards the rightness of an action, or apparently positive concepts count against performing an action (Dancy, 2004b, 108) makes his position seem both more radical than it really ought to be, and weaker than it could be. So my suggestion is for particularists to say that assisting someone isn't always good, and therefore not always helpful. It is simpler to leave prevalenced concepts out of a discussion relating to how reasons get valenced, which is the subject matter of the question "what grounds what?" in this case.

My conclusion is that, contrary to how other philosophers have understood it, particularism does not lead to the controversial conclusion that kindness can be bad, or murder can be right. My understanding is far more boring, in the sense of being conventional. Insofar as thick concepts are understood as an entanglement of a natural ground and a thin deontic or evaluative concept, I simply concede that thick concepts have invariable valences, which is the popular view in moral philosophy anyway. But I think this is irrelevant to their being put to use in the flattening debate. What we really want to know is *when* to call something kind, or to call something murder, and how particular contexts influence this.

This picture of thick concepts means that they ultimately suffer from the same problem as the default reason and descriptive generalisation views. It seems undoubtedly true to say that, in our lives in this world, thick concepts such as justice and kindness are central for us.

[3] One could actually develop this point in more detail by further distinguishing between, for example, murder and manslaughter. But this limited definition is enough to illustrate our point here.

But metaphysically, this claim reduces to the idea that certain sets of natural facts *often* ground certain reasons, which we then "package" together with evaluations, and understand as specific thick concepts. Väyrynen's "Nasty World" complaint still seems to apply, where it is possible that out of all the actions that would be called kind were they thought good, only a minority will actually be so.

2.6 Virtues

Our discussion of thick concepts leads obviously and naturally to the related, perhaps almost identical topic: that of virtues. The thick concepts properly refer to classes of actions (e.g. kind, just) whereas our virtues are the character traits, or dispositions, of the same (the kind or just person). Here we ought to remind ourselves about how virtue ethics in the history of western philosophy originates in Aristotle, who claimed that we cannot aim too much for exactness in ethics, who seemed to have a contextual understanding of ethics, and who recommended that we focus on developing just those stable character traits that will enable us to respond reliably to all kinds of situations. Perhaps virtues are a particularly friendly concept for the particularist who wants to be right about the "true nature of morality".

For these reasons, some modern philosophers (e.g. Price, 2005, Leibowitz, 2013) have suggested that Aristotle was the first particularist, or, in the other direction, that the modern position called particularism has an Aristotelian character. Indeed, the *Oxford Dictionary of Philosophy* defines moral particularism as:

> The view owing its ancestry to Aristotle, and defended by the contemporary British philosopher Jonathan Dancy, that in deciding the rights and wrongs of action general principles are of little value: the devil lies in the details." (Blackburn, 2016, 316)

Now, there are two separate things we can do here. Firstly, we can decide about this point made just now. Should we see Aristotle as a particularist, or modern particularists as Aristotelians? We can answer this question simply by looking carefully at both views, and drawing an informed and considered conclusion about how similar they really are. And tellingly, Dancy, to whom the particularist view primarily belongs, thinks not. Dancy thinks that his view, and the views of Aristotle, are importantly different in such a way that we ought not to seek to elide them, however tempting it might be to do so (Zingano, 2013).

The second thing we can do is to drop the idea of a specific alliance or comparison between Aristotle and Dancy, and instead just take the concept of a virtue directly, as analytic philosophers are wont to do anyway, and consider what use we might put it to here, if any, in our discussion of the flattening objection. The analytic moral philosopher who has written the most on the how and why of a virtuous particularism without seeking to specifically endorse an Aristotelian understanding is Constantin Sandis, who summarises his position as follows:

> A generalist virtue ethic about character traits (understood as reliable dispositions to feel, act and perceive the world in certain ways) is compatible with moral particularism concerning actions and reasons…I refer to the resulting view as Particularist Virtue Ethics. (Sandis, 2021, 206)

And why is it compatible? Firstly, Sandis draws a distinction between the reasons one does an action, and which values inform that action. It seems that, for Sandis, a strong particularist view would be one where something bad can inspire a right action. Whereas, on his view, even when a type of action usually thought to be wrong, is instead found to be right, we will only ever find goodness in the grounds. He writes that when someone tells a lie for the right reasons, they will not be acting dishonestly or be a dishonest person. This seems similar to my own understanding of thick concepts: they are only applied when they ought. We can draw three key points from his text:

1. It is often the case that an action ought to be performed that has a quality commonly taken to be important or central e.g. a kind action, an honest action, or a courageous action.
2. Such actions will be reliably performed *when they ought to be performed* by agents who have the related stable disposition of character e.g. by kind, honest, and courageous people.
3. Properly understood, it will never be the case that a virtue is what inspires a wrong action, or a vice inspires a right action. We ought to limit our application of virtue concepts in this way.

I will take these points in reverse order. On the third point, we can refer to both my earlier discussion of rule fetishism (Gleeson, 2007) and my understanding of thick concepts. The worry that a focus on virtues, and being virtuous, can actually lead us astray, is voiced by

Stangl, as well as Dancy and Christine Swanton (Stangl, 2008, Dancy, 2003, Swanton, 2001) with a form identical to that of Gleeson's rule fetishism. All three make the curiously controversial claim (as Dancy with thick concepts) that kind actions may not always be right, or relatedly, kindness may not always be right making.

A focus on being a kind person, and concomitantly, on focusing on doing actions of a certain kind, can privilege an abstract commitment over a practical response to context. However, I don't think this problem arises here in the same way as with Thakral's generics view, for example. Sandis wants to preserve the primacy of the particular situation, and wants to define the *truly* virtuous agent as one who really can tell which actions will actually be kind, and therefore ought to be performed for that reason. While the point about virtue (or rule) fetishism has *some* weight, we can define ourselves out of it here.

But this leads to my second repeated worry, which is that this limiting of Sandis' view is too tidy. We seem to need to pin the notion of a virtue to an action, e.g. the virtue of kindness to the kindness of an action, and say that an agent that possesses the first will competently perform the second. Indeed, this is how virtue ethics is sometimes understood: as either somehow being predicated ultimately on action,, or needing to be action guiding (Hursthouse, 1999, Duclos, 2023). When calling actions kind, for example, we assign these labels in a *verdictive* way. We see the action, we decide if the label "kind" is appropriate, and if it is, we apply it.

Returning to my discussion of the thick *deontic* concept of murder, to call a killing a murder is to pass verdict that it was actually wrong, but it will be wrong for some other, independent reasons. So this progression takes us back to where we started with thick concepts; no ultimate why is provided, but we seem to be left with the basic particularist insight that rightness and wrongness will really be informed just by reasons in particular situations, and only after can we absolutely, confidently judge which thick concepts, and concomitantly, which virtues, may be assigned.

But of course, the view is presumably meant to be that the truly virtuous agent will be the one who judges *in advance* that an action will really be kind, and then performs it for that reason. But the judgement here will be made for some independent reason, and therefore the kindness will not be the reason it will be performed. It is a standard move in metaethics to say that if a concept adds no extra explanation to the rightness or wrongness of an act, it cannot be the reason for doing that act. When we compare the statements "there are reasons to do that act, that make it right" and "there are reasons to do that act, and reasons to also specifically call

it kind, and therefore it should be done for the reasons it should be done, and also it would be kind" the addition of kindness doesn't clearly seem to add anything to the doing of the act. It explains why we might appreciate or promote or praise the act in a specific way, but it doesn't seem to be one of the original reasons to perform the act.

On the second point in the list above, it does seem fair to say that the truly virtuous agent will be the one who knows when one actually ought to perform an action of a particular kind. But whether this is true or not, again, we want to know is the *why* of the when the action should be performed. While it must certainly be true that virtues are an important concept in moral theorising and in moral life, our question here is a very specific one: whether virtues are useful for solving the flattening objection. And that objection is meant to be one about basic explanations of how we consistently decide which actions are right or wrong.

I suppose my position here stands outside of the generalism/particularism debate, but is really a more general one about virtue ethics: that virtues, or virtuous actions, will always in the end need to be taken together with some independent discussion of reasons about what *makes* actions right or wrong. And neither "it was done by a kind person" or "it was the sort of action which is called kind" are the right *making* facts; the basic explanatory ones. They are certainly important and relevant considerations, but they do not explain right and wrong making at the basic level.

This is how (as in my discussion of thick concepts proper) Dancy, Stangl, and Swanton get to saying that thick concepts are themselves contextual, in that kindness may not count in favour of performing an action, and may even count against it; something that Sandis denies, and which I concur with him about. Dancy wants the independent reasons to perform the action, the basic deontic facts, to have absolute primacy, and a secondary consideration that the action will be performed by a kind person, or has the apparent form of kindness, may be misleading or irrelevant. I think the correct solution is, like thick evaluative concepts, to understand virtue concepts as verdictive, which allows us to avoid making unusually controversial claims like "kindness is bad", but also makes thick or virtue concepts unamenable to addressing the flattening objection.

2.7 Subjective Contexts

I wish to briefly remind the reader that moral particularism is not only a form of nonnatural realism, but is only meaningful as such, in that it requires the existence of both natural grounds and dependent moral facts to claim that the relationship between the two is contextual. I say this because it could be tempting to attempt to respond to the flattening objection via some kind of metaethical naturalism. This is where we attempt to ground morality in the human animal in an empirical, scientifically explicable way, using data drawn from fields like evolutionary theory, or psychology. Plausibly, such data could form part of a story of why we take certain sorts of actions to be central, or why we are often inclined to act in a certain way. But, on particularism proper, this sort of story could never be *just* the explanation required; not because it is particularism, but because it is nonnaturalism. However persuasive such a story is, it will be irrelevant here.

However, there is another way we can utilise a focus on either on individual character, or generally on human nature, and that is as descriptions of the lives and contexts of either particular people, or groups of people, and concomitantly, the prescriptions they individually or jointly recognise as centrally important. This approach involves accepting the assumption that individual persons or groups of persons usually have a regular form of life, such as their career, or their role as a parent, for example. Perhaps a judge would overwhelmingly find reason *to* tell the truth, or to focus on the justness of a situation, with this being nonaccidental and nonarbitrary. While there may be no, or very few, *global* default reasons on this picture, there would be a great number of *local* reasons that really are "default" for particular groups or individuals. And since it is individuals that navigate the moral landscape, presumably understanding these defaults would be useful and meaningful.

While I think this way of looking at things *will* end up being relevant to dealing with the flattening objection in the context of *epistemology* (which I will focus on in the fourth and fifth chapters of the thesis) I think it is insufficient for our purposes here, i.e. as a *metaphysical* solution to the objection that promotes some concept of general truth. If we take again some of the specific practical concerns we looked at in chapter one: the need for moral education, the need for shared social values, the need to be able to justify ourselves to others, this subjective way of understanding context would only ever partially address these concerns. It would certainly allow us to fill out these topics locally, whereas, in this chapter, we are specifically

searching for a *solution* to the flattening objection that solves it precisely by including something *global.*

To return to my discussion of semantic generics above, I mentioned two reasons to worry about generics: seriousness, and difference in kinds. But there is a third that becomes clear when considered together with our "local moral landscape". Many of those domains which possess generic norms are precisely the ones that belong to particular ways of life. There are surely both default reasons and generics endorsed internally to the communities of soldiers, athletes, academics, doctors, and so on.

For example, in the training of judges, certain values can certainly be promoted. Internally, among the society of judges, there can be shared expectations about thought and behaviour. There can be shared standards of correctness which judges can reference in explaining or justifying their decisions to each other. All of this can be repeated for athletes, for policemen, for scholars, for soldiers, and so forth. But this is only of limited use in considering what is meant to be general, or global.

2.8 Utopianism

The final way I think we might respond to the flattening objection is by way of utopian or ideal theory. This sort of "weak particularism" doesn't strictly involve identifying general truths in the real world we live in, as the other seven examples I surveyed do, in one way or another. Rather, it involves abstract thinking about what perhaps the general ethical truths *should* be, and importantly, outside of any particular situation. I shall call this sort of musing on ethics transcendental. Various existing philosophical theories seem to refer to such an ideal. For example, on classical utilitarianism, everyone applying the utilitarian calculus perfectly appropriately and correctly would presumably create, if not a truly ideal world, at least the most ideal one possible. Political philosophies like communism or anarchism contain a picture of the ideal organisation of human society. Systems of religious ethics, like those found in Christianity or Islam, contain a clear understanding of how we should ideally behave and ideally be.

So this would involve grounding our selected central truths in a picture of a perfect world or society, as the virtue approach does in the ideal person, or the subjective approach does in the ideal of a particular form of life. This is perhaps the mirror opposite of the default

reasons approach based in our observations of our own world, or Väyrynen's theorised "nasty world" (Väyrynen, 2004). Little, in response to such a suggestion, cautions us:

> Given that our own struggles are animated most centrally by the presence of moral imperfection, it is folly to think we can best understand morality by starting with a removed ideal—say, the Kingdom of Ends—and then adding layers of failure, rather than imagining incremental improvements to the imperfections we confront. (Little, 2001a, 38)

One way I think we can understand the thought here is by using the familiar language of reasons, which both sides of the debate prefer. Presumably we would define the "Kingdom of Ends" as one where *there is no reason* to lie, steal, and so on. But there is nothing in particularism that disagrees with this. Particularism is simply the view that context *can* make a difference. It contains no specifically positive claims about when there *will* be reasons to act in discord with any moral principles, or that this will even ever happen. This makes the ideal of a Kingdom of Ends a self fulfilling prophecy. Again, it is simply built into the definition that there will be no reasons to lie there, leaving the particularist nothing to even potentially disagree with. Really, in the Kingdom of Ends, there would be no distinct "particularism" or "generalism"; since there would be nothing to disagree about, to compare or contrast, the distinction would simply collapse and lose all meaning.

Perhaps we can imagine a world where none of the unfortunately messy situations particularists dwell on will be thrown up. There would only ever be reasons to tell the truth, and only ever reasons against violence. Whether the topic is particularism or not, presumably these are uncontroversial features of any morally perfect world. In this case, the principles and the reality would be in complete agreement, and the ontological difference between particularists and generalists would dissolve.

Again, particularism has no essential notion of variability built in: it is simply the observation that such variations can happen in our world; while particularists also endorse and make much of the stronger point that such variations in fact do happen in our world, this second point is not even necessary for our point here to stand. Both particularists and generalists express their theories in terms of evaluations, and the grounds of those evaluations. In theorising the Kingdom of Ends, we are precisely deciding in advance that, regardless of how this world is, in *that* world certain evaluations really will firmly belong to certain facts.

One could ask the further (and in many ways far more substantial) question of where we would get our firm picture of a morally perfect world from anyway. The two main options on the table seem to be attempting to think normatively from some outside or higher viewpoint, or deriving and abstracting such a picture from observing our own world. Both of these are problematic in different ways, at least according to the way nonnatural realism is usually set up. But answering these questions is unnecessary for our purposes here (Although I will return to it later, in another context, in chapter five of the thesis). The flattening objection is about and inspired by our own world, and secondly, is about a dispute between different forms of nonnatural realism. So, I don't think any kind of ideal or utopian thinking will be the right way to address the objection.

2.9 Conclusion

In conclusion, a nonexhaustive but representative list of ways we might attempt to incorporate the non-flatness requirement into a particularist story is default reasons, generics, defeasible generalisations, evaluative generalism, thick concepts, virtues, subjective contexts, and ideal theory. I have attempted to show here why I think all of these approaches are in some way problematic when it comes to solving the flattening objection. In the following chapters, I discuss how these failures lead us to reconsider what our metaethical theories are really capable of saying, and describe a different way to understand particularism that seeks to dissolve the force of, rather than solve, the flattening objection.

This will involve simply accepting that particularism offers no unified, structured, explicable, *global* picture of the external moral reality, as its opponents want it to do, and its proponents here seek to do. And this will be because it is impossible for *anyone* to provide such a picture. Instead, it will end up not mattering. Particularists will not need a general, global picture of the external moral reality to refer to, this need will be misguided. And this is because particularism has a perfectly clear understanding of how we actually *make* moral decisions in any *local* context, and this was the way the flattening objection was practically motivated anyway.

3. Quietism

> ‘Quietist’ realists reject, and in doing so raise challenges to, various commonly shared assumptions in contemporary analytic meta-ethics. (Akhlaghi, 2024, 2)

In this chapter I proceed as follows. I first give an argument as to why the responses in the previous chapter both do not and cannot succeed, and I give this position the name of “quietism”. Secondly, I consider the ways this label is commonly used in the literature, and point out some relevantly important differences between my own position and that of the other so called quietists. Finally, I make some remarks on how we should proceed.

In the previous chapters I introduced what I take to be the most popular criticism of moral particularism in the recent literature: the flattening objection. I not only surveyed a range of philosophers who directly make the criticism against particularism, but also a range of philosophers who indirectly implied the objection by way of considering some specific moral phenomena, and concluded that particularism fails to explain them.

I then considered eight ways particularist philosophers have attempted to respond to the objection by way of *solving* it, which means to change the particularist theory such as that it meets the requirements the proponents of the flattening objection set out. I explained why I think each one of those eight responses *individually* fail to solve the objection, and gave some specific reasons for each, but while doing so, I introduced some more general thoughts that might apply to any such attempt.

I will now turn to argue directly that there is one common reason why not only these eight responses fail to solve the flattening objection, but why no solution to the objection can be provided. But before I explain why we cannot solve the flattening objection, let’s go over two things again. Firstly, what background considerations inspire the objection at all. And secondly, the shared background commitments generalists and particularists have in common.

3.1 Background Considerations and Commitments

And the flattening objection, I think, springs intuitively from two commitments that all the generalist philosophers I discussed share: moral realism, and metaethical correspondentism.

Moral realism is, of course, the view that morality is "real", in the sense of being mind independent. Moral truths exist "in the universe" and are discovered rather than invented, and their truth does not depend on anyone thinking them true. There is an external moral reality and to have moral knowledge is to have knowledge of this reality; to be morally correct is to be correct according to this reality. Further, some realist philosophers go so far as to think that knowledge of this external moral reality is *required* to be moral, to be morally successful, on purpose rather than by accident.

Metaethical correspondentism, to reiterate, is the view that our metaethical theories should correspond with our ordinary moral lives and all the phenomena we find there; the task of the metaethical philosopher is to explain the thing we are already doing, and a metaethical theory that seems flagrantly at odds with what we already *know* morality to be, will be incorrect. This is the spirit in which the philosophers I survey point out apparently obvious truths about our actual moral lives in their criticism of particularism.

These two views taken together lead the generalist realist philosophers, who observe that we apparently both have and use general truths in our ordinary moral lives, and who endorse the core realist tenet that ethical truths stand outside of particular minds, places, or times, to think that the external moral reality is composed of general truths, with the knowingly correct practice of morality being the application of the same.

Simply put, if both realism and correspondentism are true, then the generalities and centralities we apparently observe in ordinary moral life will exist in the external moral reality, in a structured and nonarbitrary way. Indeed, in the other direction, presumably it is the case that they are generalities and centralities in ordinary moral life *because* they are facts about the external moral reality, since realist philosophers usually think we have successful knowledge of that reality. To these philosophers, it is obviously a matter of descriptive fact that saying things that are not true pretty much always matters morally, while shoelace colour pretty much never matters.

These are not insights gleaned from descriptive ethics in the formal sense. The field of descriptive ethics involves actually researching and surveying ordinary people to discover their moral views, commitments, and practices. But Hooker does not get his claims about how moral education "normally" works by surveying 1000 parents and recording their responses. Instead what we have here are apparently commonsensical assertions which presumably these philosophers draw informally from their own experience of ordinary moral life. These philosophers do not *argue* for these claims in the way that we argue for ethical theories, for

example. The combination of realism and metaethical correspondentism is what drives their confidence in simply asserting the obvious truth of their claims about what is dominant or central in ethics, and demanding that particularists modify their theory in light of the same.

Secondly, what do generalists and particularists agree about? Well firstly, generalists and particularists agree on the above: that realism and correspondentism are true. Particularists also take themselves to be explaining what is really going on in our ordinary moral lives, and how we get to be knowingly right and wrong. Generalists and particularists also both agree about ethical grounding, which indeed, all metaethical philosophers agree about. They think that moral facts are grounded in natural facts.

They finally agree that these moral facts are understood in terms of reasons: the natural facts ground reasons for or against certain actions or conclusions. The only relevant disagreement for our purposes here, then, is that the particularist thinks the move from a locally corresponding claim (however broadly we construe this) to a global one is a step too far. On orthodox particularism, at least, however many times you find that lying was wrong *here*, it does not entitle you to claim that lying is wrong everywhere.

The stakes, as set up, are that the generalist thinks that without saying something general or global about ethics, we can neither do justice to realism or correspondentism. The practical purport here is about how and why people both can be and are regularly morally correct, and correct according to the external moral reality. This is why the generalists wanted the particularists to try to solve the flattening objection, and why the particularists concurred, and tried to do so.

3.2 Why the Flattening Objection Cannot Be Solved

Anyway, what is my general argument for why the flattening objection cannot be solved? While my explanation will rely indirectly on the discussion in the previous chapter about the failure of various particularist responses to the flattening objection, that alone does not of course form any complete proof that *any* such attempt will fail. I identified eight ways particularists have responded; might there not be a ninth? So I will now provide an argument which directly explains why we should really go on to adopt this view, and also serves to explain *why* those responses ultimately cannot work. And this will be an argument that the facts the flattening objection wants us to provide are *epistemically inaccessible*, which means we have no method of knowing them.

The point I make is not that the facts do not exist, or that facts are not important or relevant or meaningful; indeed, particularists very much do think the facts exist, and it is facts that are important or relevant and meaningful. Further, I make the point in the strongest way. I do not say that it would merely be very difficult to know the facts. I also do not say that we have merely not yet got them worked out properly, and may in the future. My conclusion will be that it is *literally impossible* for anybody, ever, to know the relevant facts to be able to *solve* the flattening objection, which its proponents want. We do not reject the non-flatness requirement because we do not want to say anything, or there is nothing to talk about, but because we *cannot* say anything. While, of course, this is a particularist view, it is not the view that *particularists* cannot say anything. It is that nobody can say anything.

So why are the facts impossible for anybody, ever, to know them? Let's return to how the flattening objection is set up, in terms of apparently acceptable answers, for a moment. The philosophers I surveyed think, simply put, that we can take moral facts and put them onto two lists. We can imagine drawing a line down the centre of a sheet of paper, and labelling the left column "central" and the right column "marginal". These two columns represent (*vis a vis* my characterisation of the flattening objection) what is common and what is rare in our moral lives; what is usual and what is unusual.

Presumably, as they would have it, "pain" can then be written down in the left hand column, and "shoelace colour" can be written in the right hand column. By considering a range of different contexts, situations, choices, decisions, we would end up filling out the left hand list. The "filling out" here is crucial. The philosophers I surveyed did not just think that we would find *some* things to put on the left hand list: they thought we would identify all the central facts, that the list would be limited. Here again is Hooker on the same "children need to learn which facts can be relevant. But they understandably hope this is some (relatively small) subset of all possible facts." (Hooker, 2000, 19).

There is also an obvious pragmatic point to be made in their favour here, irrespective of whether their view is true. The limited list of central moral truths is meant to be known by pretty much all ordinary human beings, memorised, applied, and so on. If someone were to argue that there were 10,000 central moral truths, the knowing of which would allow somebody to always make correct decisions in any and all contexts, then we could simply respond that nobody will ever learn and remember them, and that if this is really the correct explanation of what it means to be truly moral, then nobody will be. This is part of what I take Hooker and the others to be concerned about. Further, for them, the left hand list seems to be the only one

that matters, and the only one we need to really bother filling out. After all, anything from "shoelace colour" to "it is raining" to "the train is late" to "the music is loud", *ad nauseum* and endlessly, can presumably be written in the right hand column, on any view at all. Indeed, the complaint of these philosophers is that particularism would provide us with some endless list of trivial facts to consider.

So what is the problem with filling out the left hand column? One reason is that a particularist doesn't think those facts are *trivial*. They don't think facts can be called important or trivial outside of, or before considering, some particular situation. So, any particularist answer to the flattening objection would have it that all facts can potentially be in the left hand column. What we are really looking for are the ones that most commonly and overwhelmingly repeat across contexts. We want to know whether, when considering all cases, there are some central or dominant pairings of natural and nonnatural facts, arranged in some certain way.

This can be demonstrated by considering the meaning of "a context". For any sort of realist, whether particularist or generalist, the relevant meaning of a context is a collection of facts relevant to some moral decision an agent has to make, whether that is taking a concrete action, or adopting an intellectual position. Specifically, the natural facts that ground some particular set of reasons in a particular place and time. Any answer to the flattening objection would tell us the reason why particular evaluative facts attach to particular natural facts. And if we take any particular place and time, we can indeed consider the relevant facts to be limited (however long the list would actually be; presumably it would still be enormous).

So firstly, we would need to know every situation everybody has ever been in, or can be in, in our world, to find some sort of "shape", "pattern", or "structure" in the data. And this, of course, is obviously impossible. This is already bad enough, but there is another, much stronger point we ought to make here. Because secondly, we would then need to know every new fact, and therefore new context, that could be generated in the future.

What we can call contextual ethics actually consists of two separate considerations. The first one is that the way our world *already is*, and has been, is capable of endlessly generating different situations, reasons, and choices. But the second is that new natural facts, and therefore (because of grounding) completely new moral problems will be created. Obvious examples are the inventions of nuclear weapons, of AI, and of manmade climate change. What all these have in common is that completely new natural facts that did not exist previously have been generated and now have to be taken into account in our decision making.

A "solution" to the flattening objection, then, would require a human moral philosopher to know, firstly, all the facts in our world, all the possible situations they can be relevant in, and how often this will occur. This is already impossible. But further, to assign moral facts the general and global status that the flattening objection wants them to have would require the same human philosopher to know every fact that will ever exist in the future.

Note that I say any *human* philosopher. The reason this matters is that particularists do not say, and do not need to say, that they can truly prove that not a single guaranteed global or general truth exists. Just that any human agent cannot prove that one does. If God (construed as the philosopher's god) exists, and we assume that God knows every fact that has ever existed and ever will exist, every evaluation of the same, and that God knows every single situation every person will ever be in, and what the correct decision to make was, then presumably God (or if we like, "the point of view of the universe") would be able to provide a complete answer to the flattening objection. This point does not depend on whether God actually exists in any particular way or not; it just relies on the fact that humans do not have any of these characteristics, and that those characteristics what would be required.

This conclusion is perfectly sympathetic to the idea that we can identify generalities *within* any particular contexts, or range of contexts. The concept of a context can be construed more broadly or narrowly. It does not just mean one very specific situation one person has to make a very specific decision in at one moment in time. A workplace, a nation, a war, a marriage, or a vocation can all be contexts. The particularist actually agrees with the generalist every step of the way but one: we simply deny the move from identifying some pattern or important concept in one place, to claiming that it counts *globally* for all people, in every place, forever.

There is a fundamental difference between holding that justice will be a central concept in the life of a judge, or that bravery will be a central concept in the life of a soldier, and holding that these centralities will hold for all people, in all places and all times. Note that my point here should not be understood in the other direction; that justice and bravery may somehow end up being trivial or irrelevant to many people, or that these concepts can *only* be central for very specific persons. They will of course be central *in* situations pretty much all of us find ourselves in. But the point here is about moving from this to a claim of strong generality or globality. It is a step too far to claim that bravery will be a central concept in the moral lives of everybody, in every place and time, in a way that it is not when making the same claim specifically for a soldier. However much information any moral philosopher may claim to take

into account in their identification of some central fact(s), if the amount of such information *globally* is truly unlimited, as I have argued, then they in fact have *no* data and therefore *no* warrant to make such a global claim. And this is because any part of infinity, however great, is zero.

Now, I think these two points are enough to definitively establish my thesis that the flattening objection cannot be solved by any human moral philosopher. But, for completeness, there is a third, complementary point to be made here, which weighs upon something I pointed out earlier: that for generalist philosophers, the list of central or dominant truths will not only be limited, but sharply limited. Pragmatic considerations alone lead us to conclude that while the identification of five or ten central ethical truths would be practically acceptable, the identification of 500 or 1000 would tend towards absurdity. And this is what I wish to provide: an argument *ad absurdum*.

To make my point more clearly, it will be useful to take the example of an unfortunately eternal (though I suppose, understandable) trend in ethics, and one that directly weighs upon the flattening objection: the use of numbers. In all systems of religious ethics, we find numbered lists of ethical duties and considerations. Christians and Jews have the ten commandments, Muslims have the five pillars of Islam, Sikhs have the three pillars of Sikhism, and Buddhists have the eightfold path. In all cases, an attempt is made, not only to state the central, meaningful ethical truths, but to do so in a very specifically limited way, as Hooker apparently thinks we ought.

In modern secular ethics, a popular example of a structurally similar view is that of Ross, who claimed to have discovered the seven key ethical truths (Ross, 1930). I will focus on this example, since it stands within the tradition and debate we are criticising: modern analytic metaethics. Why is it absurd for Ross to be able to identify *seven* key truths, that together, cover all situations we may find ourselves in? I say it is because one could only have the standing to make this claim if one had actually really seen *every* situation a human being may ever find themselves in. But actually, any of us only ever see our own small share of such possible situations, limited by time, immanence, etc. This is what is meant by epistemic inaccessibility: we literally have no way to access the knowledge required to support the claim.

Presumably if Ross was confronted by a situation that stood outside all the examples he had considered so far, he would have to propose an eighth ethical truth to accommodate it. A ninth, tenth, etc could be added *ad nauseum*. If Ross limited his view to some particular context, for example by claiming to supply seven key considerations for a philosophy professor

living in England in the early 20th century, then a particularist would immediately find much less to criticise. The *only* illegitimate move, on particularism, is the "that's it", where the identified truths are considered to hold truly globally.

And we can only assume that Ross would only bother trying to identify and write down the seven central ethical facts if he thought that this was the right project to be working on in the first place, the sort of thing moral philosophers both ought to be doing, and can do. I think that all of the philosophers I surveyed in the first half of the thesis have a similar stance to Ross. They have an idea of what moral philosophers ought to be talking about and doing.

3.3 Why Do I Call This Quietism?

I call this view "epistemic quietism", because it requires us to be "quiet" about global metaphysical ethical generalities, but for epistemic reasons. This is a popular way in moral philosophy to describe the rejection of some particular question. That we find a question to be the wrong question, or that answering it is not necessary for our purposes, or that it cannot be answered, for example, can lead us to say that we will be *quiet* about it. Philosophers that take this sort of approach are referred to as *quietists*, and their stance is given the name of *quietism*. Now, the name "quietism" has been given to a great range of views in both the ethical and metaethical literature (indeed, in other literatures also) such as that, without specific context, it really has no particular meaning, apart from that, of course, it somehow involves being "quiet" about something.

A comparison can be made with terms like "real" or "natural" or "relative" or "objective" or "subjective" here: these terms are often used as jargon terms, as technical terms of art that have a very specific meaning in a particular debate in a particular literature. Since the reader will no doubt be familiar with at least some of the philosophies or philosophers the terms quietism and quietist have variously been used to describe, I will begin by clarifying the range of meanings they can be given, before establishing precisely how I will apply these terms.

And in the literature, the name quietism has variously been given to the philosophy of Ludwig Wittgenstein, separately the moral philosophy and the logical positivism of other philosophers that were inspired by Wittgenstein, to so called "relaxed", "nonontological", or "minimalist" realism, to antirealism, and to scepticism. Now, the way I use the term has some important similarities but also some important differences with these schools of thought.

Farbod Akhlaghi provides the most recent and apparently comprehensive survey of so called quietist philosophy, the different ways one can be a quietist, and the motivations that drive the taking of such positions. Indeed, Akhlaghi concludes that we should stop using the term quietism precisely *because* it has been used so broadly that it no longer clearly refers to any particular view, and it can be misleading if we are told some particular position is "quietist". And this is the reason I am writing this current passage: it is indeed frustrating that such a wide range of views can be pushed together under a single heading in a reductive way, and it is important to distinguish them so we can understand what is actually interesting or unique about each. But it is still true to say that what they all have in common is that they counsel us to be quiet about something.

Akhlaghi identifies the common feature the various views that metaethical philosophers call quietist have as their being *countercultural*, in that they are minority positions taken up in opposition to the orthodox way of doing metaethics, and always involve rejecting *some* popular methodology, idea, or trend in contemporary metaethics. Akhlaghi identifies nine theoretical commitments that are apparently popular in mainstream metaethics, and says that a quietist view is one that rejects at least one of these nine claims. I take six of these to be relevant for the kinds of metaethical theories I consider, and these are:

> 1. There is some distinction between ethics and meta-ethics.
> 2. There is some autonomy between ethics and meta-ethics, such that at least some views in each sub-discipline do not affect the plausibility of answers in the other.
> 3. There is some (even if elusive) distinction between ethical and non-ethical language.
> 4. Moral metaphysical questions are substantive, non-trivial questions about the nature of reality distinct from first-order ethical questions.
> 5. Non-naturalistic moral realism requires the admittance of entities into our ontology that are inconsistent with naturalism about reality.
> 6. There are notions of 'the world' or 'reality' which we can fruitfully appeal to when engaging in moral metaphysics and ask how to 'locate' morality 'in the world'. (Akhlaghi, 2024, 9)

Akhlaghi divides all of the extant apparently "quietist" philosophers into two of the groups I named earlier: Wittgensteinians, and moral realists (of which, of course, particularists are a

kind). He then further divides them based on precisely which of the one or more above commitments they reject.

I shall briefly comment on the Wittgensteinians first. Akhlaghi identifies a range of philosophers who are quiet in a very specific technical sense: they are inspired either by the logical positivism that comes from Wittgenstein's *Tractatus* (Wittgenstein, 1922), or by Wittgenstein's later thought in his *Philosophical Investigations* (Wittgenstein, 1953). I make no claim to be a Wittgensteinian, or to provide any exploration, or much in the way of explication, of these schools of philosophy here. My aim is only to point out that this seems to be a common understanding of what a philosopher is referring to when they invoke the term "quietism", and to clarify that this is not what I mean to do.

So instead, rather than focusing on quietism as belonging to a particular *tradition* following a particular *thinker* or some particular *texts*, or the thinkers who work from those texts, I will instead focus directly and generally on the concept of quietism itself; what it might mean to be quiet, what we might be quiet about, and why we might be quiet about it, as well as what follows from any such specific position.

The second family of quietist philosophers Akhlaghi identifies are what can variously be called the relaxed, minimalistic, or nonontological realists. These are apparently realist philosophers who deny (contrary to received wisdom) that realism *requires* any metaphysical part. Their justification of quietism is that we ought to be quiet about metaphysics because moral realism *does not require* an external metaphysical moral reality to be true and meaningful.

What is interesting here is that particularists and particularism find no place in Akhlaghi's survey. Particularism is actually rather mainstream and orthodox in that it shares the vast majority of popular assumptions in the field of metaethics: particularists reject *none* of the nine claims Akhlaghi suggests quietist must reject, and certainly none of the six I focus on here. So according to his analysis, quietism will not be available to particularists. But they *do* reject something, and there is something they are apparently "quiet" about.

Indeed, quietism can also be used as a pejorative: as a complaint that a philosopher or philosophy is quiet on some important question, one that actually it should have something to say about (Akhlaghi, 2024, 23). Here again is Dancy on the same:

> [Brännmark]: “One accusation that has been made against your views is that of vagueness. Would you say that part of the reason for that is that particularism is a spectrum of possibilities or is there something else?”
>
> [Dancy]: “Sometimes this accusation just means that I don't give any solid guidelines, which is true but hardly a complaint – it is just a description of my position. But sometimes people say that the view is essentially vague in places where it ought to be precise, and that is more worrying.” (Brännmark and Lind, 2008, 8)

If we agree it *ought* to be precise we are agreeing with the challenge presented by Strahovnik and the others, and therefore, would want to meet it. We are accepting the charge of quietism in the negative, pejorative sense. By rejecting the challenge as misguided, we are instead standing by our quietness, because we think we have good reason for it. And that is what I think Dancy and the other particularists ought to do instead. This is how the flattening objection and the literature about it that I examined earlier in the thesis should be understood. What has annoyed the critics of particularism is that they have somehow refused or failed to answer what is apparently an important question. But if my view is correct, *nobody* can answer the question, which means particularists cannot be specifically blamed for it. Anyway, what other ways might we positively understand quietism?

Doug Kremm and Karl Schafer, in their survey of quietist metaethical philosophy “Metaethical Quietism” characterise quietism as an *attitude* a philosopher takes towards certain philosophical questions in their work. Following Jane Friedman (2013) they go on to specifically describe quietism as a *question-directed attitude*, and provide the following list of such possible attitudes:

> Unintelligible: The questions within Q lack meaning or sense.
> Indeterminate: The answers to the questions within Q lack determinate truth-conditions.
> Presupposition Failure: The questions within Q involve a false presupposition.
> Irresolvable: The questions within Q are in principle rationally irresolvable.
> Inaccessible: The answers to the questions within Q are epistemically inaccessible.

> Irrelevant: The questions within Q are irrelevant to the underlying concerns that are supposed to make these questions interesting or significant. (Kremm and Shafer, 2017, 1)

Further, Kremm and Shafer, in common with Stelios Vervidakis, distinguish quietists by which ethical or metaethical questions they have these concerns about, noting that while a "global" quietism is possible, the vast majority of philosophers prefer to identify certain concepts, topics, or questions that they find problematic in these ways: what they include within "Q". One can take a quietist attitude towards philosophy generally, towards a particular field of philosophy, towards a particular topic within that field, or a particular question within that topic (Kremm and Shafer, 2017, Vervidakis, 2008). The terminology of "global" and "local" can perhaps be used here again. Particularists will, at most, be a very local sort of quietist.

Anyway, Kremm and Shafer, like Akhlaghi, also identify Wittgensteinians and nonontological realists as fitting within this schema, and add a further group: the moral antirealists. Antirealists, of course, reject moral metaphysical theorising because they do not think the external moral reality exists. For them, there is nothing that actually exists to describe here. This is a stronger claim that the nonontological realists, who claim to be realist but merely demur by suggesting their position does not *require* a metaphysical theory. The antirealists are certain that the subject matter of moral metaphysics actually does not exist. Vervidakis, in his "Varieties of Quietism", concurs that quietism ought to refer to a stance, an approach, or an attitude, rather than a theory proper, and that specifically problematic concepts, topics, or questions are usually identified that we ought to be "quiet" about. Vervidakis adds a final group of philosophers to the three I already identified: the traditional sceptics, in the Greek sense (Vervidakis, 2008, 158).

Now, the position I describe certainly does fit into the broader schema these philosophers sketch out. I do say that particularists should be quiet about some particular philosophical questions, and I do it for a reason Kremm and Shafer identify: epistemic inaccessibility. But otherwise, my own position not only has almost nothing in common with all of the schools of quietist philosophy these philosophers identify in their apparently comprehensive surveys, but in many ways actively stands against them. Particularist philosophers make no particular claim to be inspired by Wittgenstein or to be pursuing a Wittgensteinian programme.

Indeed, I cannot find more than a handful of isolated references to Wittgenstein in the particularist literature. I think that, at most, any similarity between the view of particularists and Wittgensteinians will be indirect or accidental. The motivations and interests of each group differ greatly. I feel this needs to be made clear, since even in the *Oxford Dictionary of Philosophy*, we can find quietism defined as "quietism in philosophy is the doctrine (associated with Wittgenstein) that there is no standpoint from which to achieve the traditional philosophical goal of a theory about some concept or another" (Blackburn, 2016, 398). The only similarity here is the conclusion that we ought to be quiet about something, and the same label being applied everywhere can perhaps obscure relevant differences.

Likewise, particularists have nothing in common with the nonontological realists, who think realism does not *require* any metaphysics. Particularists think realism does require metaphysics, and supply plenty of it. Indeed, again, Dancy defines particularism as a view *about* moral metaphysics, that helps itself to, and makes much of, uncontroversially mainstream moral metaphysical concepts like nonnatural kinds, and ethical grounding. My position here is not that we do not *require* moral metaphysics, just that we are limited in how far we can go in describing it. I call my position *epistemic* quietism because, while its target is moral metaphysics, the argument and claim are entirely epistemic. In fact, like other particularists, I make no negative claims about the *existence* of any external moral reality, just about our knowledge of it.

Finally, it goes without saying that particularists not only have nothing in common with antirealists or with sceptics, but are actively opposed to such views. Particularism is a form of moral realism that is entirely committed to the idea that in each and every case there will be a *real,* correct answer about what ought to be done. Further, particularists are actively antisceptical. They think there are plenty of moral facts, and that we know exactly what they are. Now, there is one thing that particularism has in common with all of these sorts of views, which I will bring up at the end of this chapter.

3.4 Metaethical or Metametaethical Quietism?

But first, I will briefly turn to draw a further distinction, by way of returning to the difference between ethics and metaethics, which I discussed earlier in the thesis. And the reason for this is that, as we can see, all the forms of quietism Akhlaghi, Friedman, Kremm, Shafer, and

Vervidakis identify are forms of *metaethical* quietism. This is apparently how everyone has understood the term in the literature so far. But when we say metaethical quietism, we mean the metaethical stance that we should be quiet about some or all *ethical* questions. In other words, while metaethical quietism is a second order position, its target is first order ethics.

In this thesis, however, I have entirely discussed metaethical questions. And it is some particular metaethical questions I will argue we should be quiet in response to. Properly, then, this sort of view should be called *metametaquietism*, because of instead of commenting on ethical talk, it comments on metaethical talk. To my knowledge, nobody has yet proposed this specific view in the way I will. There are philosophers who are *antimetaethical* in that they think that moral philosophy does not require metaethical theories. There are also philosophers who are generally quietist about ethics, and therefore might find any ethical question meaningless, whether first, second, or (as here) third order. But moral particularists are not antimetaethical, or generally quietist. Moral particularists are in the business of speaking confidently about metaethics.

3.5 Where Do We Go From Here?

Now, where does all this leave us? At first glance, apparently at something of an impasse. Their endorsement of moral realism and moral correspondentism leads the generalists to make a certain challenge to the particularists; and I suggest they should be quiet in response. If this is all there is to say and the debate ends here, then (at least from the generalist point of view) the charge of quietism in the pejorative sense fits, in that particularists claim an apparently important question has no answer. Those who are already wedded to particularism will presumably continue to be sympathetic to it, but anyone who thinks the flattening objection has any weight will struggle to be persuaded by what might otherwise be attractive about particularism. If we think particularism is the preferable metaethical outlook, how are we to persuade the broader field of the same, without losing what makes the view unique or interesting?

I think the way to go here relates to the positive understanding of quietism: that of quietism as being “therapeutic” (Akhlaghi, 2024, 3, Vervidakis, 2008, 162). And this is what particularists really do have in common with all the other quietist philosophers and philosophies we find in the moral literature. The meaning I give to “therapeutic” here is to persuade philosophers from looking at, thinking about, and talking about the wrong thing, to

looking at, thinking about, and talking about the right thing. And the way we can do this is by again remembering that particularists are also realists and correspondentists. Particularists do not deny that moral education, or intrapersonal justification, for example, are genuinely important moral phenomena that need to be explained, and further, they entirely agree that they ought to be explained according to a realist schema.

What we need to do is remind ourselves what moral metaphysics is meant to be *for:* why were we ever talking about it in the first place? By returning to the practical concerns that lead philosophers to formulate a non-flatness requirement in the first place, we can begin again at the start, but this time, move in a more appropriate direction. And I think these concerns can be satisfied by a clear explanation of how anyone gains moral knowledge *inside* a particular context. If we have a perfectly acceptable understanding of how people regularly know what is right and wrong, we satisfy the realist and correspondentist worries that motivate the complaints we found about education, or about society, for example.

In other words, while we are quiet about metaphysics, we can find plenty to say about epistemology. And since all decisions are made locally anyway, a sound particularist epistemology confined only to successful *local* knowledge would suffice. Moral education is the education of how to act (each time) in some particular situation. Moral justification is the justification of how we acted in some particular situation. As well as epistemic inaccessibility, Kremm and Shafer identified "irrelevance" as a characteristic of quietist views: that quietists think some particular formulation of a philosophical question is in fact irrelevant to explaining or understanding the real concerns that motivated that question, and this is just what I think we should conclude about the flattening objection.

So in conclusion, in common with every other form of quietist, I reject the challenge the generalist philosophers make: that we should "solve" the flattening objection. Instead, I think we should *dissolve* the force of the objection by finding another way of looking at things that satisfies the real concerns that drive us to such metaethical theorising in the first place. Again, particularists, like most other metaethical philosophers, are correspondentists who are inspired by, and seek to understand as accurately as possible, our ordinary moral lives. As Akhlaghi writes:

> A quietist, at first pass, aims to dissolve, as opposed to solve, one or more philosophical debates. Once we recognise that a given philosophical debate is based upon some false presupposition, then adopting any view sharing those

> assumptions to solve it would be confused. Bereft of its false presuppositions, the debate dissolves away. (2024, 2)

Kremm and Shafer also identified the charge “presupposition failure” as one characteristic of quietist philosophy. And I do think the flattening objection, which dominates the particularist/generalist debate, is based upon at least one false presupposition: that moral philosophers are able to, and ought to be in the business of, making global claims about the external moral reality. But in this case, the debate will only dissolve away once some alternative picture is put into place. And this is what we will achieve in the next two chapters.

4. Moral Knowledge as Case-Based Reasoning

> Our account of the person on whom we can rely to make sound moral judgements is not very long. Such a person is someone who gets it right case by case. To be so consistently successful, we need to have a broad range of sensitivities, so that no relevant feature escapes us, and we do not mistake its relevance either. But that is all there is to say on the matter. (Dancy 1993, 64)

The point we have reached so far in the thesis is as follows: briefly, we have identified that both particularists and generalists are united in sharing a realist view of the external moral reality, and that particularists have a persuasive explanation of why we cannot provide any complete general or global description of that reality. I have argued that this should lead us to accept a quietism about moral ontology, insofar as it is epistemically inaccessible, and to attempt to dissolve the flattening objection as a formal problem about metaphysics.

But the concerns that motivate the flattening objection are still relevant and available for discussion. Questions in moral metaphysics only have the meaning or importance they do because of humanistic concerns. And while, as I have identified, the flattening objection formally has a metaphysical form, all of the worries that inspire it have an epistemic basis: they are about how people actually make and justify decisions. If we recall again the examples I began with (such as moral education, or social order) we can see that an argument about how we decide what is right and wrong on a case-by-case basis, and therefore *locally*, ought to satisfy those concerns in practice.

Our central question seems to have become "how do we know what is right or wrong on a case-by-case basis?". The metaphor of the "moral landscape" that the flattening objection makes much of is useful here. It is uncontroversially true that any actual, concrete ethical decision is made by a particular agent in a particular context; in a particular place and time. Therefore, our worries about "shape" or "structure" in ethics ought to be *practically* satisfied by a description of our ability to consistently detect *local*, rather than *global* shape. Even if that global shape is, as I have argued, epistemically inaccessible, it is *local* shape that tells us perfectly well how agents know what is right or wrong on a case-by-case basis.

One point I wish to make at the outset is that *some* answer to our question is guaranteed pretheoretically. And this is because in ordinary life we are certain that there are good moral judges: people who consistently manage to make the right choices. Ethics clearly gets along

just fine without agents being provided with sophisticated philosophical theories. Therefore, the naïve answer to our question "how do agents consistently evaluate the situations they find themselves in?" is "they just do!". Apropos the previous chapter, I think that this point is one that all those who engage in the particularist/generalist debate will readily allow. Neither generalists nor realists are sceptical about morality, and both camps have a correspondentist methodology. Success in ordinary moral practice is precisely what is held up as evidence or example on both sides.

Further, the theory we are looking for is not one that would actually be used by an agent to make the right choice; this is not what I think the flattening objection requires at all. Our question is not one about action guidance, and has nothing to do with ethical theory proper. Any or none of the prevalent ethical theories could be true, and it would not weigh on our question here. Instead, at most, it should simply be able to explain *why what is already perfectly well happening is happening* while remaining neutral at the level of ethical theory. After all, generalism and particularism are abstract, metaethical positions anyway. While they differ on whether agents should, must, or need to use general principles or truths in their decision making, they are formally open on how *exactly* this should be specified. One can hold a deontic form of generalism or of particularism, for example, based on whether one thinks the duties are grounded locally or globally. Generalism and particularism are both standardly cast in the language of reasons for or against something, and we can say something concrete about the having of these reasons without going so far as to specify any particular first order ethical theory.

Moreover, a particularist answer must be one that relies firstly on the content of particular situations; it must be entirely "internal". This is simply implied by the core particularist tenet that one cannot look outside or above a situation for abstract, general advice. Insofar as genuine moral knowledge and competency (and therefore the achievement of the same) certainly exist, the success must be explained by information available in the particular situation, and the characteristics of the agent who is in that situation.

And that leads us to our second point, which is that there is one kind of general information we can help ourselves to: general claims about human thought, or about reasoning. This is different from the weak particularist views which I previously discarded in that those views attempt to find general truths located *independently in the external world*. While the difference may be subtle, I don't think it bears on the metaphysical problems pointed out there to make general claims about human decision making. After all, surely the whole point here is

that the thing we carry from situation to situation is *ourselves*, and it is ourselves who become wise and educated, who become good judges. General truths about human decision making, for example, will not themselves be general *ethical* truths.

Particularists describe our ability to "get things right case by case" as a skill. This skill is variously given the more specific name of judgement, discernment, competence, ability, or wisdom, and for our current purposes these terms can be used interchangeably.

4.1 Desiderata

In discussing what they say about this skill, I will keep in mind four points that I will return to when necessary, which are largely inspired by a range of conclusions or observations made earlier in the thesis. These are:

1. **The skill clearly exists since it can be descriptively observed.**

   Whatever the correct theory of particular moral knowledge is, again, it should account for the fact that there are clearly already plenty of competent moral judges; plenty of people who are morally wise. The theory is not one that *an agent needs to possess to be wise*; it is instead one that describes what is already very much happening. It will describe why wisdom, which clearly exists, is explicable without our needing to claim that any general principles or truths exist. In other words, the wise will possess the competence, not the philosophical theory, and the latter will explain the former, not the other way around.

2. **The skill is learnt over time.**

   The point here is that the explanation should account for moral progress, for moral education. We do not think that we enter the world as perfectly competent moral judges, and it would be remarkable if our ability to "get things right case by case" never improved, or was incapable of improving, or indeed, regressing; since moral degradation, or regression, is also known, and a complete theory of moral knowledge should explain all the phenomena we substantially have that relate to the same. A key part of the flattening objection was that particularism was meant to problematise how morality could even be learnt, if it only allowed for "one-off insights". Here again are Lance and Little on this complaint:

> The view seems to imply that there is no structure to moral theory at all. Moral understanding must be simply a matter of accumulating a series of one-off pieces of insight (x is good here; y is good there; do z next Thursday in Pittsburgh), a picture which makes it puzzling how morality could be learned. (Lance and Little, 2004, 436)

So one thing our explanation of local decision making should do is make it unpuzzling how people get to be good at decision making.

3. The skill extends into unusual and novel situations.

This is a key desideratum for *any* theory of moral competence, but perhaps one that a particularist theory should particularly incorporate. On generalism, we carry our limited list of moral principles around with us in our heads, so it is unmysterious how we deal with new situations, however unusual or surprising. Ultimately, we keep applying the same rules! But the particularist thesis is directly inspired by the concept of a new or unusual situation, so it would be strange if the particularist account of moral competence was mysterious regarding novel contexts. This point should be taken together with #2, as it is precisely our learning and becoming wise over time which presumably increases our chances of confronting a new and difficult situation. Unusual and novel cases are those that really highlight the difference between generalism and particularism, in terms of the *apparent* weakness of the second, as the weight of the flattening objection only really becomes clear in such circumstances. This is because in simple, obvious, and familiar cases, particularism and generalism actually just offer up the same answers. Without the concept of the surprising or unusual case, we probably wouldn't even have a moral particularism, or a flattening objection to the same.

4. The skill is relevantly similar to other skills in other normative/reasoned domains.

This is another key tenet in the particularist literature: particularists do not think that moral reasoning is different from other kinds of reasoning, or that moral knowledge is different to other kinds of knowledge. The correct theory of moral competence is really meant to be a more general theory of *normative* competence. Dancy himself prefers to describe his view as one about practical reasoning in the "philosophy of action", and

on *his* particularism, at least, what we are really talking about is a general skill of decision making, of making choices. On this view, choosing which of two countries to visit on holiday, or choosing which of two houses to buy, will involve the same sort of thought as a moral dilemma. Here is Dancy on the supposed difference between moral and nonmoral reasons:

> nobody knows how to distinguish moral from other reasons; every attempt has failed. How does that fit the suggestion that there is this deep difference between them? Not very well at all. (Dancy, 2004b, 76)

This is the epistemic picture derived from the particularist metaphysical claim that *any* piece of information can matter in a moral decision, and vice versa, in any other kind of decision. Now, this idea is not one that I personally endorse; I challenged it earlier in the thesis, for example, when I discussed why considerations about generic statements in other domains might not so easily translate to the special case of morality, and here, I will add further reasons to challenge it. I think it is misguided and ends up causing much confusion in the literature, leading particularist philosophers astray; there are, in fact, some things that are special about moral decisions that are not special about other kinds of decisions, and here, since I have specifically turned to focus on epistemology and decision making, I will say what these are. But in what follows, it will be useful for the reader to understand that the majority of the contributors to this literature *do* think in this way.

Now how does Dancy, to whom the particularist thesis primarily belongs, describe this skill? For reasons that will become clear, I initially wish to focus on Dancy's definition of this skill as one that is developed both from and over the experience of many cases, and Dancy gives us the following two passages that endorse this sort of understanding:

> What we are trying to do is to establish what reasons are present in the case before us. The ability to do this is a sophisticated one, which children develop as they grow up; presumably it is one for which some form of training is virtually essential. If we want to know what it is like to have that ability, we could start

> by asking what it is that competent judges bring to a new case. Of course, one thing that they bring is their experience of similar cases. (Dancy, 2004b, 142)

And,

> To know the practical purport of a concept is to know the sorts of difference that its applicability can make to how one should respond. But knowing this is just being able to tell the differences made case by case, in a way that is informed by one's past experience. (Dancy, 2004b, 143)

So here we have a skill that already exists in the world, that is learnt over time, that extends into novel situations, and is similar to competence or good judgement in other domains. Dancy's initial definition fits well with the four desiderata we said we would keep in mind while trying to define this skill.

4.2 Case-Based Reasoning

Other particularists agree that our "skill" in discerning local shape has such features. Uri Leibowitz provides such a picture, where there is a general definition of "skill" across domains that is grounded in experience of past cases, here using the example of motorcycling on different kinds of terrain:

> According to one particularist model a person of practical wisdom obtains moral knowledge by exercising a skill of discernment. It may, therefore, be useful to consider the scope of other skills by way of comparison. For example, consider a skilled off-road motorcyclist. Let us suppose that she acquired her skill by riding in certain kinds of terrain. If she is a skilled rider, though, we should expect that she will have no difficulty riding in a slightly different terrain from that with which she is experienced. Moreover, we may expect that she will do better than unskilled bikers in unusual and unfamiliar terrain— perhaps we can expect her learning curve to be steeper in new environments. So to the extent that we have a handle on the notion of a skill, and to the extent that we can make sense of the acquisition of moral knowledge by the application of a skill, we

> should expect this skill to extend to new and even to somewhat unusual situations. (Leibowitz, 2014, 41)

Elsewhere, we can find moral expertise explicated by way of comparison to, for example, chess expertise (Swartwood, 2020, Dancy, 2004b) football, musical interpretation, literary appreciation, architectural design, medical diagnosis (Bakhurst 2006) economics, psychology (Zamzow, 2015) car repair, and jazz appreciation (Dancy, 2010). Particularist Little concurs with the general thrust of my discussion here so far and of the above examples, claiming that, on particularism, "wisdom consists in the ability to discern and interpret the shape of situations one encounters" (Little, 2001a, 32) and "wisdom requires life experience" (Little, 2001a, 36).

Now, for reasons that will become clear, I will turn to specifically focus on a particular type of comparison between domains: one where a further comparison is made between human thought and AI. Nancy Salay provides a view that is similar to Leibowitz', in that for Salay, moral competence is gained by experience of many cases. But Salay goes on to make a further move here: not only does she compare moral reasoning to reasoning in nonmoral domains, she compares human reasoning to the "reasoning" of AI software, which she also takes to be relevantly similar:

> A good cognitive model for moral particularism, one that is also used in machine learning, is case-based reasoning (CBR). In CBR, past experiences dictate how future decisions will be made, but they do so with the help of what I'll call local generalisations or patterns, as opposed to the global ones required by generalist moral principles. And it's this distinction that really matters for Dancy… The driving thought behind CBR is that a good deal of reasoning can and is achieved by individual case comparisons, by finding one case, which was successful in a particular way, that is similar to some other case having a similar goal. (Salay, 2008, 395)

Salay uses the example of a machine which is trained to distinguish between photographs of male and female human faces, and over time achieves a higher and higher success rate at doing so. She wants to have it that in the moral realm, there is a relevantly similar skill of noting similarities and differences that is iterated over time, and importantly, one that is similar across domains (here, distinguishing faces):

> Suppose that a network is designed to distinguish between female and male faces. Initially, the system's activation values and weights are set to random so that one can expect about a 50% success rate. The system is then trained on a set of inputs and outputs; given a female face input, for example, all of the weights are slightly adjusted so that the activation value of the female face output node is slightly higher than that of the male face output node. After successive training sessions, such networks can achieve very impressive success rates. (Salay, 2008, 392)

Relying on these examples from machine learning, she goes on to conclude "we achieve the enviable state of moral competence via case-based reasoning" (Salay, 2008, 404). If this were true, it would straightforwardly be the answer to the question we have arrived at in this thesis, since we want to know how a particularist can be morally competent. Before I turn to consider firstly whether we ought to use such examples, and secondly, whether they really do support the conclusion that case-based reasoning explains moral competence sufficiently for our purposes here, I will show how Salay's view is shared by other particularist philosophers. Dancy himself also makes an attempt to define the development of human moral wisdom with the example of machine learning in his "Can a Particularist Learn the Difference Between Right and Wrong?", where he writes:

> Connectionist machines…consist of two or more levels of …nodes. There are input nodes, which 'fire' in response to …features of the…environment…There are also output nodes, whose levels of activation are determined by the degrees of activation of the nodes below them … Finally, between input and output nodes there are normally 'hidden nodes', which…facilitate the processing of information in the system. There are normally no direct linkages between input and output, that is; everything travels via the intermediate, hidden level. Such a machine consists entirely of these nodes and of the connections between them. Nodes 'fire' in response to features of the machine's experience or to the firings of other nodes…a system of this sort is capable of changing and developing over time, and so, in effect, of learning better ways of responding to input. (Dancy, 1999, 67)

Now, I will just take Dancy's word for it that a machine of this sort is capable of changing and developing over time to become better at responding to input. While he was writing in 1999 when such programmes were much less developed than they are now, at this point it has become common knowledge that such AI models improve over time, at least at the tasks they are used for (say, identifying images of leaves). What matters for our purposes is any possible similarity between such machines or programmes and human moral decision making. And Dancy goes on to explain what he thinks the relevance is:

> …we should distinguish between occurrent representations and standing representations…there is a system's representing now that this particular act is generous…this is the occurrent representation…there is the system's having a representation of generosity stored somewhere within it, in virtue of which it is able to recognise generosity…this is the standing representation…This distinction is important because it is in the latter kind of pattern that the system's…competence, is stored, while it is in the production on occasions of the former kind of pattern - patterns of activation - that standing knowledge is demonstrated. Importantly from the particularist point of view, at the level of standing representations there is what is sometimes called "deep" or "strong" context sensitivity… what pattern of connectivity represents a particular concept…varies depending on a net's training …It seems to me that a machine of this general sort should be quite capable of developing…links in response to instances of a resultant concept, even if those instances share few common features, and certainly even if there is no one pattern common to them all. Particularist accounts of the relation between wrongness and the wrong-making properties case by case would cause it no difficulty…this…would merely be the analogue of a child's learning the concept 'wrong'. (Dancy, 1999, 67-68)

What is interesting here is that Dancy seems quite sure that the learning and "thinking" of such a programme is the correct way to understand how we, as humans, make moral decisions in a particular way. He specifically invokes one of the examples we looked at all the way back in chapter one, as part of my explanation of how the flattening objection is grounded in specific

moral phenomena, and in this chapter, as part of what our definition of moral competence ought to be, which is the moral education of children. Finally, he tells us that:

> With a sufficiently wide range of initial 'given' cases, the machine should become perfectly capable of coping with the fact that there are many different ways of being wrong, and that properties that on one occasion count in favour of acting can on another count against…The machine will not only be able to respond appropriately to a fully presented situation, in which all relevant information is fairly given. Given an initial but incomplete description of a proposed action it would be able to make something of an informed guess as to whether further information would lead to an overall judgement against it or in its favour. And given defective information it would still be able to cope, in the manner of which connectionist machines are recognised to be capable. In all these ways it seems to me that the machine is a perfectly adequate model for particularist accounts of moral rationality, or of competence with the concepts of right and wrong. (Dancy, 1999, 68)

The key part of this quote is where Dancy relies on "the manner of which connectionist machines are recognised to be capable". It is certainly true that such machines are widely recognised as capable at identifying and classifying various kinds of information. But the only thing that matters is whether this would hold for *moral* information, and Dancy is unfortunately silent on this point. He simply assumes that there is a relevant similarity here. The way I prefer to understand this is by remembering that Dancy, and other particularists, think there is nothing special about moral reasons or moral thought. There are only natural facts, reasons, and practical conclusions about what actions ought or ought not to be done. I invite the reader to just keep this in mind for now.

Dancy actually describes how such "connectionist machines" work at great length, but I have tried to capture the relevant points I wish to focus on throughout this chapter and in each of these examples, which are the analogy between human thought and "machine thought", and of this being the right kind of example to prove that human moral wisdom is the product of experiencing many cases. Finally, moral philosopher Gilbert Harman is also hopeful about the prospects of using advances in machine learning to explain moral wisdom from a particularist viewpoint:

> Suppose you want to be able to classify various instances on the basis of … features like color, size, temperature, etc…you might want to classify leaves, or handwritten symbols, or viruses…A classifier is a system that takes…a cluster of observed features and outputs a classification… A perceptron is a simple classifier that has…weighted inputs, one for each of the…features an object can have.
>
> The perceptron is trained using data… If the perceptron outputs the correct classification of that datum, no change is made is made to the perceptron. If the perceptron outputs the wrong classification… the inputs are changed … This procedure is repeated going through the data many times. It can be shown that this procedure will eventually yield a perceptron that does well at classifying all the data, as long as the data can all be correctly classified by some perceptron. Once the data are correctly classified, the perceptron can be used to classify new cases. (Harman, 2005, 49)

What is essential to note here is that Harman independently arrives at the same view as Dancy and Salay in two important respects. Firstly, his examples are all of natural, empirical kinds, like leaves. Secondly, he simply assumes that this is the same way we classify moral information, without actually explaining why (I will not repeat my educated guess about why this is). The key phrase here is "as long as the data can all be correctly classified by some perceptron". Well, can *moral* information be classified by "some perceptron", like leaves can? That is precisely what is up for discussion here, and precisely what this entire line of thought depends on. And again, Harman doesn't actually tell us. Harman gives the case-based reasoning method the machine uses to make decisions, and the one we are focusing on in this chapter, the name "transduction", and concludes that:

> Transduction performs considerably better than other methods in certain difficult real-life situations…where there is relatively little data…moral transduction, if it occurs, leads to the acceptance of a specific moral conclusion that is not based on the prior acceptance of a moral principle that covers that case…We can now ask whether it actually ever occurs. How could we tell whether people ever use moral transduction or, indeed, any sort of transduction? This is a difficult issue.

> We cannot simply appeal to the fact that people seem to reach reasonable moral conclusions without being able to state moral principles. (Harman, 2005, 50)

So Harman thinks that his explanation in terms of transduction fits well with the moral phenomena particularists want to explain, and this is surely true. But it seems to involve simply assuming that particularism is true, assuming that machine thought is alike to human thought, and finally taking the explanation for granted. Harman goes on to propose and respond to potential objections about the use of such computational examples in a debate about morality:

> Objection: Statistical learning theory is concerned with reasoning to a classification of a new case given data consisting in information about the classification of other cases. But moral reasoning is not reasoning about how to classify something, it is practical reasoning about what to do.
>
> Reply: Reasoning as to whether a certain act is wrong is reasoning as to whether to classify the act as wrong. Similarly, reasoning as to whether to do a certain act is reasoning as to whether to classify the act as something to do.
>
> Objection: Moral reasoning is not reasoning from data.
>
> Reply: Sometimes in moral reasoning people compare the present case with other cases which they (think they) know how to classify as right or wrong. That counts as reasoning from data. (Harman, 2005, 52)

At least for cases that are similar to some other case, then, apparently this comparison is a sensible one (we will discuss the problem with such similarities later in this chapter). Anyway, we can see that at least some particularist philosophers favour this line of thought, which focuses on machine learning, as best explaining moral competence on particularism. And it could be said that there is a serious problem with the use of these sorts of examples in a debate about morality. References to AI are becoming more and more popular in the ethical literature, often based around questions whether AIs can have an ethical outlook, whether AIs can be used to aid the making of ethical decisions, the ethical risks for humans in using them in such a way, and so on. An entire PhD thesis could be dedicated to any one of these questions, and they are

beyond our focus here. It is simply beyond the scope of this thesis to draw a conclusion on whether the "thought" of these machines is similar to human thought in *any* relevant way for moral philosophy.

As I noted above, neither Dancy, nor Harman, nor Stangl, actually bother to answer this question. They are all confident that the comparison is apposite. Now, as with many philosophical questions, this can be related back to more basic and abstract questions in other parts of philosophy. And the question of whether machine thought is alike to human thought surely relies on substantial conclusions drawn in the philosophy of mind, which I have neither the space nor expertise to do justice to here.

So instead, I will simply assume, as these philosophers seem to do, that there is some relevant similarity here. And this will be unproblematic because the conclusions I will draw will not depend on whether this is actually true. For the point I wish to make, it will suffice to focus entirely on the "companions-in-guilt" angle being pushed in these discussions, where the way an AI (or human) reasons about a nonmoral domain, helps us explain how humans reason about moral questions. And on nonmoral domains, it seems certainly true anyway that these machines are successful in case-based reasoning in nonmoral domains. It is not controversial to point out that AI programmes have demonstrated consistent success in tasks like playing chess, or identifying photographs of dogs. The controversy revolves entirely around their inclusion in the field of ethics, and I will continue to consider ethics humanistically anyway.

Now, I wish to make the following two points about this perhaps initially persuasive picture. Firstly, it is not a basic picture of moral knowledge, and fails to tell us where we actually get knowledge from in the first place, learn it over time, and consequently, carry it forward into novel situations (as in #1 - #3 above). Secondly, the explanations of it doing so fail because they are of a companions-in-guilt style form, which rely on #4 above to get started. The reason for this is that while moral reasons are indeed similar to other reasons, those other reasons have one important difference with moral reasons.

4.3 Why Doesn't Case-Based Reasoning Explain Where We Get Knowledge From, or How We Use It in Novel Situations?

So why is the view initially persuasive, but ultimately insufficient? At first glance, these points about past cases seem to contradict the particularists ultimate theoretical claim: that there is no guarantee that one case will be like another, that a reason in one case can be an opposite reason,

or no reason at all, in other. But this is too strong when thinking *pragmatically* about our actual world. The key word here is *can*. Particularists do not claim that situations or contexts *will* have different evaluations; they only claim that they *can*. There is nothing in particularism that prevents a particular agent from being presented with relevantly similar situations again and again. Indeed, commonsensically this is what we should expect, since individual people so often have regular forms of life.

Insofar as our question is "how do *particular* agents get things right case by case?" it is perfectly reasonable to point out the obvious truth that particular agents have particularly regular forms of life, and that their past experience *will* be overwhelmingly useful going forward. Consider the life of a philosophy professor. During a long career, he regularly finds himself in situations which he finds to be relevantly morally similar. A student who is unwell asking for extra time, say, or alternatively, a student who is caught cheating. Over time, he believes that he becomes wiser, more sensitive, and more competent in dealing with the situations fairly. As he goes from home to office and lecture hall and back, there are equally situations he pretty much never finds himself in for years on end. Indeed, Cambridge moral philosopher C. D. Broad ends the introduction to his *Five Types of Ethical Theory* by admitting that:

> My range of experience, both practical and emotional, is rather exceptionally narrow even for a don. Fellows of Colleges, in Cambridge at any rate, have few temptations to heroic virtue or spectacular vice; and I could wish that the rest of mankind were as fortunately situated. (Broad, 1930, 24)

While I chose the example of a philosophy professor here, the point can perfectly well be made for any form of life. While case-based reasoning certainly explains how a person becomes skilful at confronting situations that are somehow similar to those that came before, it has nothing to say about how somebody would confront a situation that is actually novel.

This is a line of thought I previously pursued in chapter 2, on the subject of grounding central ethical truths in the regularities of particular forms of life, and here I wish to develop it by extending it specifically in terms of whatever wisdom we get from that form of life being translatable so some completely different context. So, while a philosophy professor may have developed sound practical wisdom in dealing with precisely those situations he is most commonly confronted with, were he thrown into a world war, a nuclear apocalypse, a famine,

or a prison, would his well developed practical wisdom, if defined in terms of case-based reasoning, equally extend to the surprising moral contexts and choices he could suddenly be confronted with?

The point here is that the case-based picture of wisdom relies on life going on as it did before, and on the future being similar enough to the past. In formal terms, philosophers who think moral wisdom can be explained by case-based wisdom think moral wisdom is induction. But, as I have tried to show, while an account of wisdom or competence grounded in past cases will be *pragmatically* so often true for any particular individual, and we shouldn't make a fuss about this, it fails to account for novel situations.

The second, and opposite point here is that the view is mostly silent on how we *get started*. It requires that someone needs to have experienced multiple similar cases of any particular "type" to be purposefully successful, rather than accidentally. Now, particularists do have two points about this. Firstly, that we are given our start externally, by our elders, our parents, by our society. And this is surely descriptively correct. Secondly, that we internally develop during our early experiences, and "come to see" the way things can, and often will, be. This again, seems correct.

But these thoughts merely highlight the lacuna. For our elders or our society to have a body of correct moral knowledge, human agents must have some more basic ability to acquire it in the first place, such as that anyone could have it in the first place. Secondly, we must have this ability to be able to internally iterate over our own experiences, however early we begin to do this.

So my provisional conclusion is that a case-based account of wisdom is correct, but incomplete; it is silent on how any of us have basic moral knowledge in the first place, and while it does indeed explain how we would confront situations that are truly similar to those we were in before, it is silent on how we carry it forward into surprising situations. Case-based reasoning, at most, will be *part* of, or a supplement to, a complete explanation of how, on particularism, we gain moral knowledge.

4.4 Companions-In-Guilt

I discussed the problem with using companions-in-guilt style arguments in the particularism/generalism debate in the second chapter of the thesis, there in relation to the concept of *seriousness*. I will now discuss it again specifically regarding the concept of a novel

situation. So why is the companions-in-guilt style comparison with other normative domains, which is unfortunately the preferred approach among particularist philosophers due to their basic commitment to *any* information being potentially morally relevant, misleading on the point of novel situations? I think it is because those other domains are locally "closed" in a way ethics is not. Even if ethical *reasoning* is structurally similar in *some* respects to say, reasoning in chess, or in football, the relevant grounds in those domains, however broad, do have an edge. I will explain this point with just two examples: chess, and the references to AI by Dancy, Harman, and Salay.

Dancy uses the chess example thus: "the competent chess player [does not need] to be aware of all the indefinitely ramifying contributions of the different aspects of the position in front of her in order to reach a responsible judgement about what move there is most reason to make." (Dancy, 2004b, 142). The key word here is indefinitely, which is incorrect, and this mistake is what makes comparisons with other apparently similar domains useless for our purposes. Because the number of possible moves in chess is merely very large, rather than indefinite. Indeed, mathematician Claude Shannon calculated the number to be 10^120 (Shannon, 1950). This is what we mean by a domain being "closed", compared to "open".

In morality, not only can any current fact ground a moral reason, but, as I have pointed out throughout the thesis, entirely new facts will be generated in the future. Whereas in chess, no new facts will ever be added, nor can *any* current fact ground a reason: only facts internal to chess can. The metaphor here would be that for reasoning in chess to be *exactly like* ethical reasoning, instead of merely *similar,* it would always need to be possible that, for example, the shape or size of the board might change, that new pieces might be added, or that the pieces might be allowed to move differently. Chess can be (and has been) "solved" by a computer, in that all the potential moves from any particular position can be calculated. Morality cannot be "solved" in this way. We can continue this thought by returning to Salay's facial recognition example:

> Suppose that a network is designed to distinguish between female and male faces. Initially, the system's activation values and weights are set to random so that one can expect about a 50% success rate. The system is then trained on a set of inputs and outputs; given a female face input, for example, all of the weights are slightly adjusted so that the activation value of the female face output node is slightly higher than that of the male face output node. After successive training

> sessions, such networks can achieve very impressive success rates. (Salay, 2008, 392)

The problem here is that the amount of female and male faces is again merely very large, rather than indefinite. The data set is closed, which is why it can be "solved" by the computer. If we were to give the computer images of *every single* face that currently exists (around 8 billion) it would achieve a 100% success rate at distinguishing them: there would be nothing more to learn or do. The amount of humans faces that have ever existed, or will exist, is of course much larger, but again, not infinite. Scholars estimate that, depending on which start date is chosen, perhaps 50 billion humans have ever existed (Westing, 1981). Moreover, there are biological, evolutionary limits on how much faces can really differ. Such a machine could be trained on *any* "closed" set and end up achieving the same results. So, if this example is meant to be apposite for moral wisdom, the set of information required for moral wisdom ought to be similarly "closed". And as I have argued exhaustively throughout this thesis: it is open.

One further question, relating specifically to our wonder about surprising situations grounded merely in the way the world *already is*, would be something like whether such a trained machine could be given a different data set and therefore immediately go on to distinguish, say, different types of stone. My understanding is that it would not. Case-based reasoning clearly holds only in reference to some sequence of genuinely *similar* cases.

The second and stronger question, which instead focuses on my definition of context as including the generation and addition of actually new information to evaluate, is whether the machine would be able to reliably parse a dataset where completely new information of constantly varying kinds is being added. Again, I think the answer is that it would not. And all of these points can be translated back into points about human thought anyway. Case-based reasoning, when justified by examples from nonmoral domains that have these two features (the data set being both closed and static) fails to explain why it means for human beings to have basic moral knowledge on a particularist account. Novelty does not come only in degrees, but in categories. There is a dividing line between "slightly similar" and "completely new". We wanted to know about what it means to be competent in terms of navigating an inexorably open set of information that is *also* open in terms of how much it can differ, not a closed set of information that we know will not importantly change in advance. Here, I am inclined to again refer to, perhaps ironically, Dancy's own view about companions-in-guilt arguments in ethics, given in another place:

> I just don't think that much is gained by finding a view similar to one's own in a cognate area which has had a good run for its money there. This might achieve respectability for one's own view (were this in danger of being denied) but could hardly count as evidence for its truth. (Dancy, 1992, 94)

Anyway, how do we get any basic knowledge before we become wise, or in a new environment? Where did anyone ever get justified basic knowledge from in the first place, such as that our wisdom, the wisdom of others, or the wisdom of our society, is not ultimately a happy accident? I think this is our real question, and it is one that unfortunately cannot be solved at the most basic level by reference to case-based reasoning, however much particularist philosophers might prefer to do so. While case-based reasoning is surely a relevant part of any complete picture of moral knowledge, it doesn't tell us where we get moral knowledge from, or how we apply our knowledge in novel situations. Case-based reasoning is all about learning over time, but whatever we are learning here, it is *not* the "skill" we are looking for here, when understood as moral wisdom, but rather, at most, a complementary additional part of the same.

5. Basic Moral Knowledge as Perceptions and Weights

> The basic, ground-level moral facts are facts about what is a moral reason for what …The basic reason-facts which we are to come to know are particular; their purview is initially restricted to the particular case. We need to be able to come to know these non-general facts, or to acquire justified beliefs about them; and our knowledge of them will be our basic normative knowledge. (Dancy, 2004b, 141)

In the previous chapter I took up the question of how agents come to know what is right or wrong in any particular case, without reference to any general principles. I considered one explanation, case-based reasoning, which proposes that a key part of this ability is the intellectual comparison of the present case to others, real or hypothetical, and that this ability is one refined by training and experience. I concluded that while this is surely part of a complete picture of a morally competent individual, of course it does not tell us how we come to have any moral knowledge in the first place, so that we can first evaluate cases to be able to compare them, or evaluate completely novel cases. So what we want to know to fill out our account is how we evaluate the particular case here *only* on its own merits.

5.1 Three-Layer Model, Grounding, Buck-Passing, Valences and Weights

To better understand what we mean by our *most basic* knowledge: the knowledge that is "moral knowledge" proper, and explain why we have any at all, let's look again at how particularists model moral knowledge. Particularists use a three-layer model of deontic concepts (Dancy, 2004b, Berker, 2007) which can be subdivided into a view about grounding (Roberts, 2018), and a view about buck-passing (Dancy, 2000).

Grounding, of course, is the view that normative facts are grounded in natural facts. The buck-passing view is the view that rightness and wrongness are not independent moral properties, or indeed properties at all. They are entirely reducible to there being reasons for the same: to say that an action is right or wrong is not to assign a property to it, but instead just to

say that there are reasons for or against doing it. In the language of realism, there is a "real" or "robust" property of being a reason, and to call something right or wrong is a semantic shorthand that actually *refers* to such reason properties, rather than to a rightness or wrongness property proper. Formally, the buck-passing view involves two claims, one derived from the other:

1. For an action to be right or wrong is *just* for there to be a reason or reasons for or against doing that action.
2. Therefore, rightness and wrongness themselves cannot give us reasons to act.

Like the distinction between natural and normative kinds, this one will also be relevant later. In summary, there are the natural facts that ground reasons for or against deontic conclusions (that an action is right or wrong), those reasons, and those conclusions. To know that an action *here* is right is to have reasons that confirm that, and to have those reasons requires there to be facts that make it so. As well as being so valenced, these reasons are held to have *weights*, to be stronger or weaker, which explains why a particular conclusion is ultimately the most correct one.

For example, if there are strong reasons in favour of doing an action, and only weak reasons against doing it, those weak reasons are held to be "defeated" by the stronger ones. If there are only weak reasons in favour of some particular course of action, but strong reasons against, the opposite holds. Moreover, when choosing between several apparently promising courses of action, the reasons with otherwise the same valence can be weighed against each other. E.g. action a has reasons that make it right, but action b has weightier reasons that make it right, making action 2 preferable to action a without ruling out action b.

5.2 Three Questions

All of this tells us that our *basic* moral knowledge *here*, which is given and explained entirely by the situation we find ourselves in *here*, and therefore satisfies the questions we were apparently left with, itself requires the answering of the following:

1. How do we know that any particular natural fact grounds any reason *at all*?

2. How do we know whether such an identified reason has a positive or negative valence, whether it counts for or against some course of action?
3. How do we know the weights of the positive or negative reasons we have identified; which ones are stronger and which are weaker?

The answers to these three questions ought to combine to explain how we end up deciding what is ultimately right or wrong *here*, and therefore what is to be done. This is also how Dancy parses the problem, since the "three layer" model makes this the obvious way to break down our questions about moral knowledge, and specifically, how we should understand our "skill" of deciding what is right and wrong in a novel case:

> Looking now not so much at what we bring to a new case as at what we need to determine about that case, we can say that we need two interlocking skills. The first of these is the ability to discern which features are playing the favouring and disfavouring roles here…The second is the ability to put this first sort of knowledge to use in determining how the various contributors combine here to give an overall answer to the practical question what is the thing to do. Though there are two distinct skills here… I would say that we have both or neither. The ability to discern the 'rational weight' of individual features of the situation would not be comprehensible in a creature incapable of using the resulting knowledge in practical decision. And a creature capable of determining what is overall right and wrong without any sort of sensitivity to the reasons at play is equally incomprehensible. (Dancy, 2004b, 143)

Answering the three questions above would presumably tell us why anyone has any moral knowledge in the first place, and how they would evaluate a novel situation which cannot simply be compared to other similar situations they successfully evaluated in the past. The first and second questions above relate to the first transition in the "three layer" model: the move from knowing facts to having reasons. The third question relates to the second transition: the move from having reasons to drawing a conclusion and therefore, to taking a course of action.

The discussion in the literature about the first transition, about knowing reasons and valences, revolves around the concept of *perception*, while the discussion about the second and third layers, the transition between them, and how the valences and weights of the reasons

combine to recommend an overall conclusion, revolves around the concept of *additivity*. In this chapter, I will take these questions about the transitions between these three layers in order.

5.3 Perception

So on the first two questions, the ones about basic knowledge and having reasons, which moral philosophers have so far preferred to discuss in terms of perception, we can find four sorts of claims in the literature: that particularists in fact *do* use the concept of perception to explain basic knowledge, that particularists *should* use the concept of perception for the same, that particularists *must* use the concept of perception to explain particular knowledge, or even that the concept of perception can be used to *support* the particularist thesis that there are no moral rules. When considering the short quotes below, we should bear in mind what the standard principled, generalist view is: and on a generalist view, there is apparently no mystery about the move from the natural layer to the reason layer, and therefore apparently no problem.

On generalism, I have a limited set of principles in my mind that I carry about everywhere with me, that together cover all possible moral decisions. No matter what sort of unusual or novel situation I am presented with, I will be able to apply at least one of my principles. And the way I know basic reasons each and every time is by selecting whichever natural features are available for evaluation (since only some natural features are morally relevant and available for evaluation) and applying the matching principle (McKeever and Ridge, 2006).

An example of a principle is "that an action would cause pain is always a reason against doing it". Since pain is a natural feature, whenever I come across actual or potential pain, I now immediately know that I have a reason against causing, increasing, or perpetuating the pain (and of course, where relevant, a reason *for* alleviating the pain). Particularism lacks both of these features: firstly, any feature can be morally relevant, and secondly, it can have either a positive or negative contribution. And I can only get the relevance or the valence by in some way *looking at* what is the case here. Jeffery Smith tells us that:

> Moral particularists eschew a deliberative and epistemic role for general moral principles in favour of a process whereby particular moral judgements are ultimately based upon a perception of the unique arrangement of moral and nonmoral facts of a specific situation. (Smith, 2002, 269)

While Benedict Smith tells us that “particularists typically appeal to a theory of moral knowledge which operates with a 'perceptual' metaphor” (Smith, 2006, 11). Note that this first definition is worded in terms of our perceiving an arrangement, while the second is worded in terms of a perceptual metaphor: we shall return to this later, but for now, suffice to say that these definitions are informal and open enough that they simply must be true. In ordinary language, to perceive just means to see or know that something is the case.

Do agents “see” that something is right or wrong? I assume that all moral philosophers would agree that according to the usual meaning, of course they do. Can the concept of perception be used as a “metaphor” for gaining moral knowledge? Again, I assume all moral philosophers would be happy with such a metaphor, precisely because any more particular philosophical view can be put into it. What we want to know more precisely then, is what specific understanding of “perceive” the moral philosophers in the particularist/generalist debate seem to have. And two pages later, Benedict Smith tells us that “particularist's attempt to give an account of moral knowledge according to a perceptual model” (Smith, 2006, 13).

This seems to be the stronger view: that they in fact formally endorse a perceptual model. Another two pages later, Benedict Smith adds the concept of a perceptual *capacity* to those of a metaphor and a model:

> One prominent strand in particularist writing involves the appeal to moral perception in an account of moral knowledge. Denying a necessary role for principles in moral reasoning can naturally lead to an emphasis upon the perception of the discernment of particulars. The metaphor of moral vision has been variously used to try and capture the sense in which moral knowledge and agency rests upon episodes of successful seeings: seeings that here are a reason to x…Thus knowledge that here is a reason to act can be achieved only by putting oneself in contact with it. Modelling moral knowledge on some form of perceptual capacity is, then, perhaps especially suited to particularism. (Smith, 2006, 15)

Here we find that using the concept of perception is especially suited to particularism, and therefore, I suppose, is what they ought to be using. Onora O’Neill provides a different sort of view: rather than claiming that particularists actually endorse perception, she seems to think they simply must, by derivation and elimination. O’Neill thinks that if any general knowledge

from outside a situation in unavailable, all that remains in terms of possible knowledge is "mere, sheer perception" (O'Neill, 1996, 86).

Walter Sinnott-Armstrong comes to a similar conclusion, here phrased in terms of intuition instead of perception: "without [principles], all a particularist has to offer is bare intuition" (Sinnott-Armstrong, 1999, 11). McKeever and Ridge concur that particularists prefer a perceptual view:

> Particularists emphasise the apparently perceptual phenomenology of moral judgement… [on particularism] moral judgement seems much more like perceptual judgement in which one 'just sees' that this case has certain features. (McKeever and Ridge, 2005, 101)

What I find most interesting about this particular reading is that we are told one would *just* see what is the case. This is one of the first times one of these philosophers seems to distinguish a particularist view from any other moral philosophical view. On generalism, for example, one precisely cannot *just* see. One first sees, then one matches what one has seen with some moral principle, which one has previously learnt and memorised. So this idea of *just* seeing seems not only specifically relevant to particularism (where no general information is brought from outside) but also relevant to our question here, which is what particularist do or must think *basic* knowledge is. "Just seeing" is at least a specific, if obviously unsatisfying answer, and one we will return to.

Further, it is of an accord with the way Sinnott-Armstrong and O'Neill choose to understand particularism: once anything general is taken away, all we are left with is mere perception, bare intuition, or some mysterious "just seeing". Here we have four philosophers, then, who prefer to understand the link between particularism and perception in the third possible way I mooted: that particularists *must* see things this way. None of these philosophers claim that they *do* or *should* see it this way. And *must* is the strongest of the possible options, in a grimly critical sense; the others allow particularists a choice of what view they have, but here they are apparently left with no options and nothing to say.

McKeever and Ridge go on to make a further, and in this case perhaps more hopeful, point; that perception can actually be taken as an argument *for* establishing the particularist thesis that there are no general principles. This is meant to be because accepting moral perception makes general principles epiphenomenal: if perception of particular situations

reliably provides successful moral knowledge anyway, what need do we have of a further principle?

If our successful understanding of what is right or wrong is entirely explained each and every time by our *perceiving* what is so, then adding the further piece of information that this also accorded with what we perceived somewhere else does not add anything to the explanation (McKeever and Ridge, 2005, 101-103). What is interesting about this move is that it sets up another theoretical option in the opposite direction: instead of predicating perception on particularism, it predicated particularism on perception. Instead of particularists becoming perceptualists, here perceptualists perhaps should or must become particularists. If we find something strong to say about perception, perhaps this will also allow us to say something stronger about particularism, too.

So again, we have four positions here: that particularist do, should, or must use perception, or in the other direction, that perceptualists should be particularists. While this claim is interesting, it is quite orthogonal in relation to the first three; since there are various ways of establishing particularism, and the claim that if perception is correct, so is particularism, is unhelpful in a discussion of whether perception is the right way of looking at things in the first place. Furthermore, McKeever and Ridge only briefly suggest that perception *can* support particularism; they don't insist on it. And it seems to me that there are ways perception can be made amenable to rules or principles also. So I will discuss whether particularists do endorse perception, and then turn to whether they should or must.

But before I do so, I will clarify exactly what is meant by perception here. Firstly, according to the ordinary meaning of perception, it is going to be trivially true according to the everyday definition of the word that if we successfully judge *here* that x is right or wrong *here*, and do it entirely by *looking* at what is going on *here*, then we have "perceived" what was right or wrong. It is a common locution to say "I saw that what she did was wrong" or "I can see that this is the right course of action", in the same way we might say "I see that it is red", or "it looks heavy".

If this was all that was meant, we could simply agree that not only particularists are perceptualists, but likely, everyone is: these ways of talking can accommodate principles. Going forward, I will refer to this as a "loose" understanding of perception, insofar as any more specific view can be readily fitted into it, and this sort of loose understanding is what I take those (such as Smith) who refer to particularists as having a perceptual *metaphor* to have. But

"perception" has a special, narrower, and more substantial meaning, both in moral philosophy, and in philosophy more generally.

I take ordinary sense perception first, as the standard move in moral philosophy is to define ordinary sense perception, then develop moral perception as somehow structurally similar. In epistemology, ordinary, empirical sense perception is defined as the properties of external objects producing and/or sustaining in us a *phenomenal representation* of that object as having those properties, via our five senses: seeing, hearing, feeling, smelling, and tasting.

Robert Audi is the contemporary philosopher who has done the most to develop not only this picture of ordinary perception, but also (as we shall see) a correlating picture of *moral* perception grounded in it. His definition of ordinary perception is as follows:

> What causes our perceiving a substance is its instantiating some suitable set of properties, commonly including at least one observable property, such as a colour or shape property. Unless the perception is strictly momentary, there will be both event causation and sustaining causation, even for my seeing my hand steadily, there is an event of my *starting* to see it, followed by my continuing to see it. That state is sustained by the diachronic instantiations of the relevant manual properties; those instantiations sustaining my seeing it is a causal relation. To see an object, then, is to see some suitable subset of properties of it, which is a matter of an appropriate causal relation between its instantiating such properties and our phenomenal awareness of them: that is, a phenomenal instantiation of certain phenomenal properties – a kind of representation – of the visible properties in question. (Audi, 2010, 86-87)

So, we have ordinary, empirical sense perception according to this definition. I will not attempt to interrogate the details of this (no doubt an epistemologist both could and would), because my focus here is on how moral philosophers like to attach definitions of moral perception to these definitions of ordinary perceptions. Further, this way of understanding perception is both typical of ordinary, empirical epistemology and typical of the debate and field we are in. Audi talks of there being *properties* in the external world that can be perceived by faculties we possess, and in realist moral philosophy, we find a similar focus on external properties and our gaining knowledge of the same. It seems to be definitionally true that on realism, if there are external properties and we have knowledge of them, then *something* like a perceptual picture

must be true. It is the obvious way to understand our knowing of properties in the world. So, since Audi's is a perfectly representative example in the literature, and he has a lot to say about the move into morality, I shall use him.

Anyway, it is worth briefly returning to our definition of nonnaturalism here, since particularism is a specific form of nonnaturalism. Nonnaturalists contrast their moral kinds with "natural" kinds, and define the latter not usually by their content (although this can also be done) but by how they are known. A standard definition is that natural kinds are those apprehended by our five senses, and are therefore investigated by the natural sciences via the empirical method (Moore, 1903, Shafer-Landau, 2003). This popular way of relating moral perception to ordinary perception, then, maps exactly onto the popular way of dividing objects and facts. In other words, it is the epistemological corollary of the standard metaphysical picture. Therefore, our question about moral perception can be translated, and related back to more general discussions in metaethics about the knowing of moral kinds.

So what is the move from ordinary perception of nonmoral kinds, to moral perception of moral kinds? While historically moral perception has, in one way or another, been characterised as a sort of "sixth sense", which detects moral kinds just as the ordinary five senses detect empirical kinds, moral perception is instead standardly conceptualised as a sort of "five and a half" sense in modern metaethics (Dancy, 2010). This is because of familiar complaints about George Moore's view about perception, which *was* such a sixth sense view, and because of the now universally accepted concept of supervenience, or grounding.

Moore held that we directly perceive goodness and badness in the world via an evaluative faculty. There was a 1:1 relation where the evaluative faculty is only and specifically for sensing the evaluative kinds, and the evaluative kinds were only and specifically appreciated by the evaluative faculty. Again, it is standard to define natural kinds *in terms of* their being detected by ordinary, empirical sense perception. For Moore, the answer to how we know basic moral facts was surely structurally similar: if we use our ordinary senses to detect natural kinds, then we must have a nonnatural sense for the sensing of nonnatural kinds (Moore, 1903).

Half a century later, Richard Hare (1952) formulated the now common principle of grounding, or supervenience: that evaluations are grounded in natural kinds, and there can be no evaluative difference without a natural difference. Again, this principle is a central feature of the particularist (and indeed, any) understanding of moral ontology and epistemology (Roberts, 2018).

The tension between these two is as follows: if the evaluation is definitely such as it is *because* of how the object is naturally, then why do the natural grounds have no place in the causal chain? The evaluation itself is irrelevant to *explaining* itself; its status is explained entirely by reference to the natural grounds. In the other direction, if the evaluation *entirely* causes our understanding of its status, then the natural grounds become epiphenomenal, or irrelevant.

This is how we come to the buck-passing view mentioned above. Since the model of natural grounds and reasons for evaluative or deontic conclusions completely *explains* both the goodness or rightness of something, and our knowledge of the same, while declining to posit unnecessary extra properties, explanatory power and parsimony are maximised, which are key desiderata for any metaethical theory. For these reasons, grounding requires us to posit some sort of "five-and-a-half" view, where the five empirical senses have an essential place in the causal chain, and the moral faculty posited somehow *overlays* or comes *together with* those senses (Audi, 2010, Dancy, 2010).

Crucially, the five-and-a-half sense view is identical to the sixth-sense view in terms of requiring some specific moral sense or faculty, insofar as it is a view about nonnaturalism: only in naturalism can we fully reduce moral knowledge to ordinary knowledge, and have a "five senses" view, and therefore, no need for any further perceptual faculty. On naturalism, feeling pain, for example, *just is* moral knowledge: nothing further is required. This is what we mean by a five senses view. One point to bear in mind here is that any definition of "moral" perception will include "ordinary" perception packaged inside it: nobody questions whether we have sense perception. And this is Audi's "five and a half" view:

> [I] grant that the base properties for moral properties are natural and have causal power if any properties do. The non-causal element in moral perception and in knowledge acquired through moral perception belongs to conceptual capacities… the experience in virtue of which a moral perception counts as a perception is causally grounded in perception of natural properties…My conclusion… is that although moral properties are apparently non natural properties, they are constitutively anchored in natural properties, in an intimate way such that seeing or otherwise perceiving the natural properties or relations that are their base suffices, given an appropriate phenomenal response, to make it reasonable to describe certain experiences as perceptions of such moral

properties… perception embodies a kind of moral knowledge. (Audi, 2010, 92-93)

So, this is the "standard" picture of what we mean *philosophically* by moral *perception,* and this is what I take those who say that particularists have a perceptual *model*, or specifically that they require *moral perception*, to mean. This is just the way that the moral philosophers I survey understand moral perceptions, as the combination of a nonnatural property, and some human capacity to detect the same. I now turn to whether particularists should, do, or must endorse perception in this strict sense.

5.4 Moral Sense and Reason Sense

Firstly, should they support *moral* perception? Well, particularists cannot use Audi's five-and-a-half view, or any relevantly similar view that speaks of perceiving rightness and wrongness. This is because natural kinds are usually defined *in terms of* how they are known, and that moral perception includes ordinary perception. If we follow this train of thought, we will naturally want to define the added moral kinds also *in terms of* how they are known. Moreover, if it is with ordinary sense we ordinarily perceive, it seems that for particularists to be perceptualists, we ought to want a moral sense.

Now, of course this was Moore's view, which I mentioned and discarded above; but the reasons for discarding Moore's moral sense only go for the sensing of the thin, verdictive deontic properties. The combination of buck-passing and grounding doesn't do away with senses entirely; it just displaces them from rightness and wrongness properties to reason properties. The grounding part of the equation is what moves us from a "sixth sense" view to a "five-and-a-half" sense view.

The concomitant buck-passing about rightness and wrongness means this "five-and-a-half" view can only be about reasons. So, loosely speaking, the answer to our question of how we know reasons is that we ought to have *reason* perception, and a *reason* sense. *Not* a moral sense. Not just particularism, but any view that incorporates buck-passing rules out *moral* perception as perception of thin properties. This is the eventual conclusion of Smith, for similar reasons:

> Particularist epistemology must be a very peculiar form of reason perception… [this is a] recurring worry about the alleged epistemological implausibility of particularism: the idea that agents can know about moral reasons by just looking at them when they are there. (Smith, 2006, 14-18)

But what is a reason sense? Just as a discussion of moral perception seems to turn into one about moral sense, and a discussion of ordinary perception into one of ordinary sense, we can turn away from wondering whether particularists ought to be interested in reason perception to what they think a reason sense is. Now, do particularists even think there is such a sense? It is surprising that, as I showed above, there is a consensus that they do, since Dancy, to whom the particularist thesis canonically belongs, entirely rejects both versions; of course that we perceive rightness and wrongness, but also, reasons for the same: "We have no moral sense, and no reason sense either" (Dancy, 2004b, 146).

His explanation of this is that since the usual picture of moral or reason *senses*, as a continuation of ordinary sense perception, require one to actually physically sense the relevant grounds. This would rule out the possibility of evaluating past cases, or hypothetical ones: which is a practice we intuitively ought to want to preserve. Now, this only rules out a reason sense as one *always* tied together with a concept of reason *perception,* and therefore a strict perceptual view. So on our question of whether particularists do endorse any form of moral perception, the answer is clearly no. But can we give some other definition of a reason sense that includes our interest in past cases and hypotheticals? Dancy, as is often the case, demurs on offering any substantial positive view. As with his particularism (there are no moral principles) and its related metaphysics (there is no limited theory that explains rightness and wrongness) his epistemological conclusion is a negative one: "Our sense that something is a reason is absolutely underivative and immediate, and can be given no independent theoretical support." (Dancy, 2004b, 160-161). Dancy gives a practical illustration of this thought with the following personal example about raising his own son:

> I was amusing myself when he was about six years old by asking him why I should not send him to bed. His answer was, 'I don't want to go to bed'. So I asked him why I should not make him do what he didn't want to do. He said, 'It would make me unhappy'. I asked him why I shouldn't do what would make him unhappy. And his reply was, 'There are no more answers, Daddy'. This

> seemed to me to be the right reply. We can often enough explain why a consideration is a reason, but eventually that desire for explanation will hit a brick wall, in the form of a point at which no more can be said, without that doing anything to destabilize the reason we were trying to explain. (Dancy, 2018, 1)

And this is explained by Dancy (in not much detail) by pointing out that such independent theoretical support could only take the form of:

> A reason why whatever is a reason is a reason: a sort of meta-reason. Particularism, at least in my hands, has seen no need to suppose that any such meta-reasons are there to be found…in my view no such theory is in fact available. (Dancy, 2004b, 160-161)

So, on whether particularists should or must have a theory of perception and sense, Dancy's answer also apparently seems to be a clear no. My own position is sympathetic to the concerns that make critics of particularism think that they should or must do so, and without disagreeing with Dancy's conclusions, I would seek to defuse the specific worries about perception and reason senses in the following way, at least in terms of them being a problem for particularism. If they are not specially problems for particularism, then Dancy's rather closed responses seem less annoying.

My first step is to simply agree that we should and must have a reason sense, but since particularists think there is nothing special about moral reasons, and that their definitions of "reason" and "reasoning" are otherwise orthodox, this reason sense should just be given the ordinary name of rationality. As rational creatures we detect reasons here as we detect them anywhere, and the allowance that we have reasons is a universal one in philosophy, as well as a practical fact: again, no one denies that we have moral knowledge, or reasons for the same.

At our most trenchant, we could simply declare that the question of whether we have reasons is a question for everyone everywhere if it is for someone, and that the mystery simply doesn't matter. A perhaps useful comparison with the origin of language: no one knows why we have language, and the debate is a famously obscure and controversial one in academic circles. But our agnosia about the origin of language is no practical barrier; since we simply

know we have language, the mystery engenders scepticism about *whether* we can speak correctly in no one.

What is the specific problem with simply declaring this? Why is simply saying that as rational creatures we can find reasons *here* variously implausible, mere, peculiar, bare, and so on? My diagnosis is that the problem is not actually particularisms rejection of principles, but its antitranscendental yet nonnatural realist character. This would mean that the mystery of how we have reasons is a problem for any philosophical position which has these two features.

Why is this the case? Let's take the nonnatural part first. Nonnaturalism is of course, a form of realism. There are two other sorts of realism available, supernaturalism and naturalism, as well as an entire family of quite different views: the anti-real ones. I will take two representative examples from the antireal family first: emotivism and relativism.

Emotivism is the view that morality can be entirely explained by emotions: to call something generally right is to say that it generally makes people happy; to call something generally wrong is to say that is generally makes people sad. On emotivism, our identification of positive and negative reasons is simply *the same thing* as our feeling positively or negatively towards it. No other moral foundations are held to be required, or to exist. If we endorse emotivism, there is no question about our identification of reasons (Ayer, 1936).

Relativism is the view that moral truths are specific to particular societies, and all that is needed for something to be right or wrong, true or false, is for our own society to hold it to be so (Harman, 1975). On relativism, our identification of reasons is simply given by our society. If our society approves of an action, we therefore have reasons in favour of it, and if our society disapproves of an action, we therefore have reasons against it. Again, no *further* foundations are given or required. Neither the emotivist, relativist, or indeed any sort of antirealist *needs* to explain where we get reasons from in quite the way particularists, as nonnaturalists, are bound to. Insofar as reasons are *grounded* somewhere, particularists want them to be grounded in an external, nonnatural moral reality, which, as Mackie (1977) famously argued, can be seen as controversially mysterious. Emotions and cultures, on the other hand, just obviously exist. Internally to emotivism and relativism, a complete explanation has been provided.

Now back to realism. On naturalism, our identification of reasons is a product of, for example, evolution. We have simply evolved to approve of certain actions and disapprove of others (Richards, 2017). To return to the obvious pain example, we have evolved to find pleasure positive, and pain negative. On the more specific combination of hedonism and

naturalism, we can even more strongly say that our reasons are simply given by pleasure and pain. It is just our nature to approve of pleasure, and disapprove of pain, and that's it. Again, no further explanation is needed; there is no mystery.

On supernaturalism, our moral reasons come from God, or otherwise from a higher supernatural reality. Whatever way we agree that we can identify what God wants (for example, by studying scripture) once again, we now just do have reasons to be for or against particular actions. A satisfactory answer has been given.

But on nonnaturalism, such moral properties need to be got from somewhere, and if they are not got from transcendental abstract thought or contemplation (which is where, for example, I understand Plato and Kant to get them from) then they can only be got from experience. So on generalism, which is also a form of nonnaturalism, there are only two places we can get principles from, by particular experience, or by abstract thinking outside any particular situation. And only the second of these is any use in our current discussion, which as I said earlier, will be predicated on the getting and having of knowledge only in some particular situation or situations.

To say that we get it from somewhere outside, divorced from specific experiences (however construed) would not be a particularist answer to the flattening objection, but instead, an abandoning of the particularism. While this is surely one way to go, it is useless for our purposes. If principles are gained from experience (as some commentators would have it, from seeing features always count the same way in many cases) generalists too need some form of reason sense to collect this data. If a generalist is also an antitranscendental nonnaturalist, as in my description, they have an equal problem here.

This is what leads to talk of either moral senses or reason senses being all that is left, or my suggestion that simply claiming that by being rational we just do have reasons being unsatisfactory. On supernaturalism we have a soul. On naturalism, at least the evolutionary sort, we have evolutionary psychology. On emotivism and relativism we have our emotions, or our culture. On transcendental nonnaturalism, like with Kant or Plato, we have a higher moral reality that we can contemplate. So particularisms problem is not the particularism: it is the combination of nonnatural realism and antitranscendentalism. In this thesis I have several times pointed out that particularism cannot be understood naturalistically, and highlighted the distinctively nonnatural character of the position, with all the specific difficulties that come with it. Dancy concurs, writing if "If one is a naturalist, one cannot be a particularist, while if

one is a non-naturalist, one inherits all the metaphysical difficulties of that position." (Dancy, 1999, 65).

If experience is the only place knowledge can be gotten, and our ordinary senses only give us ordinary knowledge, some reason sense is required here. And generalism is also a nonnatural view, with its only putative advantage on this apparent mystery being principles. But if principles are argued independently to not exist in the way generalists would like to have it, they fare no better on this score in the end. Recall again that Smith (2006) worries that particularism will be "epistemologically implausible" because it *must* conclude that we can know reasons just by looking at them. But one common way of understanding generalism in terms of how we actually get moral principles is that we derive them from observing cases. In this case, generalists also see reasons "just by looking at them", they simply make the further move of claiming to notice a pattern in the information they get in this way. So perhaps we can simply declare that we have reasons.

And on this point, it is important to remember here that our general position, and the debate we are within, is an antisceptical one: generalists and particularists are united in thinking that we do have reasons, and we do have knowledge. So the unsatisfactory flavour of our answer doesn't make us lose confidence in the reality of what we wish to understand; our concerns relating to the flattening objection are ultimately practical ones.

So if we allow on both sides that we do have reasons, then perhaps considering the *weights* of these reasons will be enough to satisfy our practical worries about how we draw justifiable conclusions about what is the case *here*. In the previous chapter, we concluded that our skill or ability at moral reasoning is one that can surely be refined through experience, but that the concept of experience doesn't tell us what the skill actually is. So my hypothesis will be that a relevant part of the definition of this skill will be that it is one of *weighing*, and that part of the definition of a morally wise individual will not just be that they can tell that some piece of information matters, but *how much* it matters.

5.5 Weighing Reasons

When I introduced and explained the flattening objection all the way back at the beginning of the thesis, I pointed out that many philosophers have complained that particularists have no concept of what it means for a feature to be central or marginal, to be dominant or trivial. Part

of moral wisdom and moral competence, it is argued, is not just knowing that some piece of information matters, but knowing *how much* it matters. And how much considerations matter is what we are about to investigate absolutely directly. Weight is importance, and if reasons in some particular case are all we have, then the central and dominant reasons will presumably be the ones that *weigh* most heavily. If we have a clear picture of this, we have a practical solution for particularists to deploy against the flattening objection, which is their biggest problem.

So what do particularists, or anybody, think the relevant details of reason weighing are? A common, simple and uncontroversial example used to illustrate weighing is coming across a child drowning in shallow water, which is safe for an adult to walk in to. The cost of walking into the water (getting one's feet wet) is held to be very small compared to the benefit of saving a life (Singer, 1972, 231). Here, then, there is a very light reason against going into the water, and a very heavy reason in favour of going into the water.

The concept of reason weight is by no means particular to particularism, but is a general one: on any moral view, presumably some considerations are stronger than others, and specifically, this is exactly what the case is on generalism. Generalist views usually allow that there can be some exceptions to the rules, but the exceptions have to be suitable weighty, and do not compromise the strength of the initial principle. E.g., take the two following examples:

> 1. My action would involve lying, and since there is a rule that I shouldn't tell lies, this is a strong reason against lying. However, if I tell the truth, I will get in a small amount of trouble: this is a fairly trivial reason to lie, so here I should tell the truth.
> 2. My action would involve lying, and since there is a rule that I shouldn't tell lies, this is a strong reason against lying. However, if I lie, I will save a life. This is an even stronger reason in favour of lying, so here I should lie.

While I will return to the specific extra details particularists add to weighing later, at this point what we need to know is that the debate in the literature is mainly over whether particularists think weights are *additive.* Thinking that reason weights can be added up is held to be essential to understanding competent moral reasoning: particularists deny this. So if we want to add weighing to our definition of the morally competent individual, we need to say something useful about additivity that allays the critics concerns. What do we mean by this?

5.6 Are Reason Weights Additive?

Additivity is the view that commonsensically follows from the remarks on the previous pages. As the name implies, it is the view that we "add up" the reasons for or against various actions, and choose to actually carry out the action that has the highest positive "score". Indeed, some central views in moral philosophy seem to both assume and require this understanding of reasoning. Descending to the level of ethical theory, and taking hedonic act utilitarianism as our paradigm example, ethics *just is* counting up the amount of pleasure and pain in a situation and adding and/or subtracting the values to decide what ought to be done.

While on such first order ethical theories the assignation of actual concrete numbers is possible (on utilitarianism, we can literally count the amount of people an action will affect, as in trolley problems) it is more usual in the metaethical literature to refer to reason weights in terms of them having *metaphorical* values, for example, John Broome, using the example of a pair of physical weighing scales, tells us that:

> These reasons are analogous to the objects in the lefthand and right-hand pans of the scales. Each reason is associated with a metaphorical weight. This weight need not be anything so precise as a number; it may be an entity of some vaguer sort. The reasons for you to F and those for you not to F are aggregated or weighed together in some way. The aggregate is some function of the weights of the individual reasons. The function may not be simply additive, as it is in the mechanical case. It may be a complicated function, and the specific nature of the reasons may influence it. Finally, the aggregate comes out in favour of your F-ing, and that is why you ought to F. (Broome, 2004, 37)

This point about the "function" perhaps not being *simply* additive is made by various contributors to the debate, e.g. Selim Berker (2007). It has also been called multiplicative, or by even more complex terms usually drawn from mathematics. But it is orthogonal to our discussion here. Whether additive, multiplicative, or some even more sophisticated term, the relevant point is that in some way the weights can be calculated or computed against one another, and that is what we shall examine. Additivity is by far the most common name given to any position that has this feature, so for simplicity, I shall continue to call it so.

On the drowning child example mentioned above, we might suggest that getting one's shoes wet has a negative value of -1, while saving the child's life has a positive value of +1000, giving an overall value of +999 in favour of performing the action. We call these values *metaphorical* because, of course, we don't claim either that we can actually compute an exact score for each reason, or that they can be represented even roughly by actual numbers. The real point is that we have a clear idea of how strong or weak each reason is, such as that we assign each some value *here* in our minds eye; when I write "1000", it should be taken to mean "very large", and when I write "1", it should be taken to mean "very small". This is the additive model.

A common example used to illustrate the way we calculate the best course of action in this way is the practice of making a list of "pros" and "cons" for and against a particular course of action. Horty states that the first written description of the additive model for reasoning towards a practical conclusion was given by US president Benjamin Franklin in a 1772 letter, which he reproduces:

> My way is to divide half a sheet of paper by a line into two columns; writing over the one Pro, and over the other Con. Then, during the three or four days consideration, I put down under the different heads short hints of the different motives, that at different times occur to me, for or against the measure. When I have thus got them all together in one view, I endeavor to estimate their respective weights; and where I find two, one on each side, that seem equal, I strike them both out. If I find a reason pro equal to some two reasons con, I strike out the three. If I judge some two reasons con, equal to three reasons pro, I strike out the five; and thus proceeding I find at length where the balance lies. (Horty, 2007, 1)

Berker provides a similar example, with an explanation of what we mean by calling this weighing process "metaphorical":

> Suppose Andy is trying to decide which of two apartments to rent. How might he go about making that decision? One natural suggestion is that Andy should do something like the following. On a piece of paper he should make four columns and, in the first column, write down the positive elements, or 'pros',

> associated with taking the first apartment; in the second column, write down the negative elements, or 'cons', associated with taking the first apartment; in the third column, write down the pros of taking the second apartment; and, in the fourth column, write down the cons of taking the second apartment. Next, he should decide the weight of each item in the four columns—that is, decide how heavily each pro or con will affect his ultimate decision. Finally, he should survey the relative weights of the various pros and cons and come to a final decision about which option has the most favorable balance of considerations in its favor. Note that in this process Andy need not be able to represent the weight of each pro or con with anything as precise as a numerical value. Note, also, that the final determination of which option is the weightiest need not involve anything as mechanical as adding up the weights of a given option's pros and subtracting the weights of its cons in order to determine a total weight that can be compared with the total weight of the other option. But despite these deviations from a strict analogy with a weighing of masses on a scale, there still seems to be a useful sense in which we can say, metaphorically at least, that Andy is weighing his options—or, as it is commonly put, that he is weighing the reasons for and against each alternative to see where the balance lies. (Berker, 2007, 113-114)

Berker calls this picture the "generalised weighing model" and suggests that it is both the obvious and orthodox way to understand reasoning to a practical conclusion. Shelly Kagan calls the same conception of listing and weighing the "governing function" and again suggests that it is a common understanding of weighing in moral philosophy.

> Let us call the function (whatever it is) that is taken to determine the overall status of the act on the basis of the values of the factors the governing function. In principle, even theories that agree about the list of relevant factors might disagree about what the governing function is like, and thus disagree about how the factors combine and interact in determining the act's overall moral status… [some philosophers] presuppose a particular view of the nature of the governing function…the function that determines the overall status of the act given of the particular factors is an additive one. That is, the status is the net balance or sum

> which is the result of adding up the positive and negative effects of the individual factors…The overall status is the sum of these and negative contributions. (Kagan, 1988, 14)

While Cullity calls it the simple picture:

> We can start with the following Simple Picture. … A fact is a reason for action when it bears a special, normative relation to it – the relation of normative support, or "counting in favour" … This relation is a vector, which is to say that it has a direction and a magnitude… There are only two directions a reason for action can have: either for or against. But the magnitude has no such restriction; it can vary just as much as physical weights can. The weights of the reasons for and against an action can be summed, and this determines how strongly supported the action is, overall, by the reasons that bear on it, for and against. …You ought to perform the action with the greatest net balance of reasons in favour over reasons against…Rational thought and action then correctly reflects this balance. Rational deliberation is weighing reasons correctly. Acting rationally is doing what you have sufficient reasons to do. (Cullity, 2018, 423)

Again, I shall continue to simply refer to it as the additive model, as that will be the relevant point. For our purposes, the relevant features of the additive model are:

1. Each individual reason has a positive or negative "metaphorical weight" according to its valence and strength.

And:

2. These positive and negative weights combine to recommend a practical conclusion.

Now, if this standard model were correct, it seems as if it would perfectly well answer our particular question. A key feature of our understanding of moral wisdom in terms of deciding what is best to do *here* will be that we accurately weigh up the reasons in our heads, then add them up or otherwise compute them. But Dancy thinks particularism is incompatible with the additive model.

Why is this? At first glance, there seems to be no reason why particularism would *require* us to be nonadditivists. The particularist thesis is simply that there are no exceptionless, general moral principles, because a feature that in one situation grounds a reason for an action, can ground a reason against, or no reason at all, in another situation. What here stops us from thinking that in each case we (in whatever more precise way we construe the point) "add up" the reasons for and against? In fact, it seems that the particularist, who focuses entirely on deciding what is most important *here*, what there is most reason to do *here*, might be particularly interested in the additive way of thinking.

But Dancy explains that the additive model requires the truth of atomism about reasons. And since particularism proper is predicated on the view that atomism about reasons is false, the additive model cannot be correct; at the very least, on pain of incoherence, a particularist cannot be an additivist. Holism about reasons is the claim that a piece of information that can give us a reason *for* a particular course of action in one situation, can give no reason at all, or reason for an opposite course of action, in another place. For example, the fact that my action might cause someone pain can recommend that action here, and no action at all, or a completely different one, somewhere else. Holism is the partner of particularism: holism being the claim about reasons and rationality, and particularism being the application of holism onto the topic of moral principles. Dancy sets out his view on this point in two key passages in *Ethics Without Principles*:

> Further, we must recognize that we are not dealing here always with a balancing judgement of the comparative weights of the reasons for and against. Common though it is to read of the balance of the reasons, and of weight, the matter should not be understood on the model of the kitchen scales. That model is far too atomistic to fit the subject matter of ethics (or of practical reason in general). On the kitchen scale model, each consideration has a practical weight, which it keeps irrespective of what it is combined with—just as a kilogram of butter weighs a kilogram whatever it is added to. A more holistic picture—one that recommends itself to me—has it that the presence of one feature can affect the weight of another. What is more, the two features need not be features of the same action. That one alternative has a certain feature may affect the favouring powers of a feature of another alternative. Judgement may need to pay attention to this sort of thing (which does not mean that no judgements are simple). And then there

> will be comparative questions to be faced at the end of the process, for once we have understood the claims of each action, in the context of whatever alternatives are available, we then have to decide which claim is the strongest—and this should not be understood on the model of kitchen scales either. (Dancy, 2004b, 105-106)

And:

> Kitchen-scales conceptions of rationality, or of 'rational weight', make two claims: 1. The weight of a rational element or reason is not affected by the weight, nor indeed even by the presence, of any other rational element. 2. Once one has assessed the separate weight of each element, evaluative judgement consists of adding up the pros and cons to see which side is weightier. These claims express a sort of atomism in the theory of reasons—an atomistic conception of rationality. (Dancy, 2004b, 190-191)

This is because adding up the weights of the reasons requires those reasons to have individual values that would be preserved independently, in isolation. The point here is simple enough. The adding up of the reasons on the additive model is the *second* stage. The first stage is deciding the weights of the reasons. If the first stage cannot be completed because of holism, then the second stage can never begin.

Dancy does not deny that we weigh reasons or that we draw a conclusion based on the weights; just that we cannot know any individual weight alone. His argument is therefore a modestly epistemic one; just as with my own quietist argument about the metaphysics about rightness and wrongness, he does not deny that the reasons actually have weights, or that we know them; only that we know them *individually*. Again, Dancy's picture is an apparently vague and negative one; he does not offer any precise model in place of the additive one, but just tells us that the additive model is incorrect, while otherwise agreeing that we decide, in *some way*, where the balance of reasons lies. My task in the rest of this chapter is to work this out more precisely.

Kagan, in his "The Additive Fallacy" has developed the problem with atomism and additivity even more exhaustively, using the concepts of *contrast*, *ubiquity*, *transport*, and *independence*, to highlight various problematic features of atomism. While he himself is not a particularist, and he is silent on whether holism or atomism is correct, this paper is of central

importance to the combination of additivism and particularism since he firmly concludes that additivity about reasons requires atomism about reasons.

Kagan uses the term "contrast cases" to describe pairs of cases which are identical apart from a single feature, and our identical or differing evaluations of them. If we evaluate the cases the same, he suggests, it is because we think that particular feature was not relevant to the conclusion, whereas if we evaluate them differently, we think it proves that the feature in question was relevant to the conclusion (Kagan, 1988, 5-12). Kagan thinks that contrast requires *ubiquity* (his name for atomism):

> The strategy relies on an underlying assumption concerning the role of factors [features] - an assumption that is questionable and should probably be rejected. Yet without this assumption, contrast arguments cannot be used to derive the conclusions we draw from them. The contrast strategy clearly assumes that if a factor has genuine moral relevance, then for any pair of cases, where the given factor varies while others are held constant, the cases in that pair will differ in moral status. … in principle any one pair of cases should be sufficient to settle the question of whether the given factor is of intrinsic moral relevance…if variation in a given factor makes a difference anywhere, it makes a difference everywhere. Let us call this the ubiquity thesis. Contrast arguments presuppose the ubiquity thesis. (Kagan, 1988, 12).

This is one half of the now familiar concept of atomism: that a feature retains its *status as a reason* and its *valence* wherever it goes. Kagan also makes a further claim synonymous with atomism, in terms of "transport":

> Most cases are complex enough that, even if we can agree on the list of morally relevant factors, it is difficult to know how to proceed so as to arrive at and defend a judgment concerning the overall moral status of the act in question. But if the additive assumption is correct, a method naturally suggests itself, for we can isolate the individual contributions of particular factors. By constructing simple enough cases, we can isolate particular features and note the strength of the reason that a given feature grounds. We can then "transport" this information back into complex and more controversial cases: observing the presence of the

> same feature, we can infer that it grounds a reason of the same strength. (Kagan, 1988, 24).

So also, atomism fixes the "strength" and therefore the *weight* of a reason, thus making it easily available to being "added up". Fourthly, he adds the final feature of atomism that sets it at odds with Dancy's holism: that the presence of other reason giving features, or secondary features, does not affect the contribution of our isolated feature.

> In discussing the implications of the additive assumption, I have done more than assume that the overall status of the act is the sum of the separate contributions of the individual factors. I have also presupposed that the size of a given factor's contribution is determined solely by the value of that factor: variations in the given factor will affect the nature of its contribution; but variations in the other factors will not affect the contribution made by the factor in question. That is, I have assumed that the size of a factor's contribution is independent of everything other than the value of the factor itself. (Kagan, 1988, 16-17).

Kagan calls this *independence*, and of course, it conflicts with a fundamental tenet of holism and particularism: that the relevance or weight of a feature can be modified by the presence of another feature, or by secondary factors. Note that Dancy allows that "Judgement may need to pay attention to this sort of thing (which does not mean that no judgements are simple)" (Dancy, 2004b, 106) on the additive model. We *may* need to pay attention to his worries about additivity, and only concerning judgements that aren't simple. This is because some situations really are so simple that they can be authentically modelled with the additive model.

Taking again our example about the drowning child. If we agree that the only two features that seem relevant are firstly that we would save a life, and secondly that we would slightly inconvenience ourselves, there is no problem simply counting one as a small negative, and one as a great positive: the features in this example do seem to preserve both their valence and weight when considered independently, in isolation. The specific problem, then, is that the additive model is not suitable for more complicated and confusing cases, which are precisely the ones that motivate holism and particularism in the first place.

If every moral situation were as simple as our hypothetical drowning child example, then presumably we would have no problem with having general principles like "always save

lives!" and hence, no need for particularism. While the ways in which a situation can be more complicated or confusing than this one are precisely open ended, here are just two examples, one of Dancy's and one of mine, to help make the point. Dancy suggests that sometimes a feature that seems positive *here* in isolation can nevertheless actually count as a negative when considered holistically, and vice versa. His own example is, rather whimsically, as follows:

> Contributory reasons are officially reasons capable of doing what they do either alone or in combination with others. But they can combine in peculiar and irregular ways, as we will see. There is no guarantee that the case for doing an action, already made to some extent by the presence of one reason, will be improved by adding a second reason to it. Reasons are like rats, at least to the extent that two rats that are supposedly on the same side may in fact turn and fight among themselves; similarly, the addition of the second reason may make things worse rather than better. Remember the joke about a New York restaurant: there are two things wrong with this restaurant—the food is terrible and the portions are too small. (Dancy, 2004b, 15-16)

Kagan instead suggests that a feature can exist that counts *both* for and against here, as do McNaughton and Rawling (McNaughton and Rawling, 2013, 245). While, unlike Dancy, he doesn't provide a specific concrete example, these are not hard to find.

Imagine that my dog is old and sick, and his condition has deteriorated to the point that only two options are immediately available, and that I am required to choose one in short order. Firstly, he can have a long, expensive, and painful operation that has no guarantee of actually curing him; he may die anyway, after even more suffering caused by the operation. Secondly, he can immediately be peacefully put down. In this situation, the fact that I love my dog can be turned this way and that to count both for and against both options. My loving him may make me consider that the operation is worth the uncertainty and expense if he may live longer. It may also make me consider that a dignified death that will definitely be peaceful and painless is the right choice.

A further point can be drawn out of these thoughts, which is one about the way moral philosophy itself is usually practiced by contemporary analytic philosophers. The drowning child example itself is taken from a popular text by Peter Singer. Trolley problems abound in the literature. In the particularism/generalism debate, it is common to present a simple choice,

and ask whether a moral principle should be honoured or ignored. All of this relies on the habit of philosophers to set up simple cases in thought experiments to use as "intuition pumps" for moral conclusions. The relevant factors are selected in advance precisely to allow for, and facilitate, the drawing of intuitive conclusions. But I think this practice is deeply misguided.

Real situations are just not this simple. Real situations can so often involve paying attention to just more than one or two pieces of information. If morality really did usually involve such simple choices with obvious answers, generalism, additivism and the rest would also be obvious. But our entire debate here is ultimately motivated by the fact that real life is not like this. Anyway, particularists apparently ought to reject additivism because of considerations arising out of the traditional atomism/holism debate, but Dancy otherwise demurs on providing a specific alternative picture. Now I will turn to the apparent problem with rejecting additivity, which, in the absence of an alternative, leave us at something of an impasse.

Berker, probably the main defender of the combination of atomism and additivism, and critic of particularist alternatives, considers the position of Dancy, and, indirectly, of Kagan, and concludes that it is untenable, and untenable because it renders the practical meaning of "reason for" or "reason against" incomprehensible. He begins by suggesting two possible precise definitions of what it means to be a "reason for action", and claims that on either, additivity is required for their coherence. The first definition is the "right-making conception", where a reason is what *makes* an action right (Berker, 2007, 127). The second definition is the "favouring conception", where a reason is what recommends a particular action (Berker, 2007, 133). In simple terms, the difference here is the difference between a metaphysical and an epistemic definition. The first one tells us what really *makes* an action right, the second one refers to our apparently having knowledge of the same.

Crucially, Dancy agrees that reasons for action have both of these features. I will reproduce here a quote I provided all the way back in chapter one, from an interview where Dancy explained his moral particularism to two philosophers:

> [Dancy]: …I now take the view that there is a distinction between being a reason why one shouldn't do something and being a reason for believing that one ought not to do it. The latter topic is investigated in moral epistemology. "How do you know that it is wrong? What reason do you have to believe that it is wrong?" The former investigation is moral theory proper, where you investigate what

> makes an act right or wrong, which isn't an epistemic matter at all. If you ask which of those two are primary, it is clear to me that it is what it is for a feature to make an action wrong or what goes to make an action wrong, and the epistemic question is secondary. So, I think of the matter as what makes the wrongness, that's a metaphysical question, and how we track this is the epistemic question, the answer to which should largely be informed by the answer to the first question. (Brännmark and Lind, 2008, 11)

Anyway, Berker claims that either of these definitions require either atomism or additivism; he doesn't insist on both, only one or the other. He thinks this because of an analytic truth about the definitions of reason *for* and reason *against.* Berker's axiomatic truth here is that if particularists want to use the term reason for, they must allow that the individual feature *increases* the amount an action is favoured, however metaphorically. Since it is the topic at hand, and since atomism has been separately argued against, I will focus on his point about additivity.

Berker claims that for a reason to contribute for or against, it must have, all things considered, a particular weight that it contributes for or against, again, however metaphorically: it must have *something* it itself contributes. At the absolute minimum, particularists need to assign particular weights to be allowed to hold their overall position. Since the atomism part is nonnegotiable, we might want to solve this by trying to reconfigure particularism to in some minimal way seem amenable to additivity.

But since Kagan and Dancy amply demonstrate that additivity *requires* atomism (Unlike Berker, who insists only on one or the other), the full result is that particularism itself becomes untenable. To successfully answer our question, which was how we come to know what is ultimately right or wrong *here* based only on what is the case *here*, requires an alternative positive picture of how we judge reason weights. Berker thinks that having "a coherent notion of a reason" requires one to be able to assign a weight to that reason. A definition of weight that allows one to assign a contributory weight to a reason without atomism would solve Berker's problem, and apparently, ours.

5.7 Another Way

I think the answer is that all we need to know about the weights of reasons *here* is their *comparative* weighting, and we get this by comparing them to one another. Consider the following analogy. I see a group of three men of different heights (say 180, 185, and 190cm). I know *immediately* which one is the shortest, which one is the tallest, and which one is in the middle. I need not measure any of the men individually to know this: I know it entirely by comparison. To extend the metaphor, knowing which reason is the weightiest is the same as knowing who is tallest. I don't need or want any ability to "weigh" it to be able to know so. It seems to me that this simple point renders the distinction between additivity and nonadditivity irrelevant. Knowing the weights of reasons or clusters of reasons by contrast liberates me from needing to precisely add them up.

This point exactly continues the spirit of earlier points made throughout the thesis, and the spirit of Dancy's position. I said previously that we need no descriptive theory of the external moral reality to practically decide what to do: knowing what seems to be the case *here* is enough each time. Here I say we need no theory of reasons having specific weights, or knowledge of those specific weights, to practically decide what to do: knowing what seems heaviest *here* is enough to satisfy the spirit, if justifiably not the letter, of the non-flatness requirement.

Compare my point about tallness here with the drowning child example above. Just as we do not need to be able to say the tallest man is 190cm to say that he is the tallest *here*, we do not need to be able to say saving the child weighs "+1000" to say that it weighs heaviest *here*. I will extend my initial discussion of what it means for a weight to be "metaphorical". We understand the "+1000" and "-1" metaphorically precisely because we know them by comparison. The metaphor was that the "+1000" really means "very large" to me, and the "-1" means "very small" to me. But the way I know this is just because I am there in the situation, and I compare them to each other.

I just *see* that one is very small in relation to the other, and one very large in relation to the other, and I know where to go. If I see a tall mountain next to a small hill, and not knowing the actual height of either, "metaphorically" assign them sizes the sizes of tall and low in my head, it is precisely because I saw them *together* that I do so. This is the strongest version of the point. The particularist can respond that it is actually impossible for a reason to have any weight without company to compare it with, and define weight entirely as a comparative

concept. This view could actually be extended to all the ways we use the word "weight". If we call a person heavy, it can only be by reference to their being heavier than other people.

On particularism, we firstly do not need to be able to evaluate a feature "atomistically", in isolation, to understand it, but secondly, we can perfectly well problematise whether this is even possible. We do not need to know anything general at all. Knowledge that a feature is heaviest *here* is complete moral knowledge. Knowing that it is heaviest anywhere else, or generally heavy, on the basis of having a certain known weight, is pragmatically irrelevant and therefore epiphenomenal, on top of being metaphysically implausible.

I think my view is a natural extension of Dancy's own view about additivity, while avoiding the additive/nonadditive debate, since on Dancy's view a feature has no weight in isolation, but only in company. Using my example of tallness, this would at first glance yield the apparently absurd assumption that the men have no height individually, but only in company. But my view isn't bothered by this conclusion, since it only *requires* them to have comparative size in company anyway. Just as I initially considered "solving" the flattening objection before I preferred an approach that "dissolved" it, here again I consider solving the apparent problem of additivity, before preferring to dissolve it.

Conclusion

This thesis was for anyone who, like myself, is sympathetic to the basic particularist starting point: that context is important, and that morality might not best be described in terms of absolute general principles, and wonders whether there might be another way to go. And as I mentioned in the introduction, and tried to amply demonstrate throughout the first half of the thesis, this starting point is clearly not taken to be obvious or sensible by a great many moral philosophers. But there were many things we found particularists and their generalist opponents did have in common philosophically, and we did find another shared starting point. And this was what philosophers call metaethical correspondentism: the idea that we take our ordinary moral lives and all the phenomena we find there as our starting point and our data. Both particularists and generalists are clearly united in wanting to understand phenomena like moral education, or social organisation, or intrapersonal justification. This is the neutral middle ground they meet on. And even though the particularist ultimately ends up going in a very different direction, they do allow that the concerns that motivate the flattening objection are real and important.

Moreover, particularists and generalists are alike in many other ways that the apparent division in the literature, I think, for example where particularism is described as being especially extreme or radical, tends to obscure. Both are actually very ordinary and orthodox sorts of nonnatural, metaethical philosophers. As well as both being correspondentist, both endorse the usual form of moral realism, both endorse the usual understanding of ethical grounding, both think we should explain morality in terms of mind independent explanatory reasons, and so on.

These are what I take to be the reasons why we get the sort of responses I surveyed in the second chapter of the thesis, where particularists attempt to develop "weak" particularist views, precisely because they largely share the theoretical outlook and practical concerns of their opponents. My suggestion has been that they should, in the end, actually stand their ground here, but for reasons also that are really at home in the usual way of doing metaethics. I don't think my own arguments or conclusions are so extreme or radical, either.

While in the literature the particularist picture of how we ought to understand moral knowledge have largely been driven by an enthusiasm about case-based reasoning, which is an approach that focuses on the concept of *similarity*, my own answer that I have ended up with is instead one grounded especially in the concept of *difference*. The difference between how

much things matter. I think both together form an acceptably complete picture, that is available to particularists when explaining their own commitment, and capable of allaying the worries and fears of the generalists.

Another commitment that these two groups of philosophers share is their antiscepticism. Both particularists and generalists think there is "real" moral knowledge to be had, and that we have it. And my own conclusion is very much in this spirit. I am not sceptical that we have reasons, or that we will know the difference between a big or small one when we see it. The flattening objection was set up in a way where it demanded that we acknowledge the difference between two sorts of moral information, the central and the marginal. And while I rejected this challenge when understood metaphysically, I think my own epistemic answer pretty much maps onto this division. That on a case-by-case basis, we will know what is central by way of knowing what we have the strongest reason to do. And that is all the particularist needs to say.

Bibliography

Akhlaghi, F. *Meta-Ethical Quietism? Wittgenstein, Relaxed realism, and Countercultures in Meta-Ethics.* In *Wittgenstein and Contemporary Moral Philosophy*. [London: Routledge, 2024]

Audi, R. *Moral Perception and Moral Knowledge.* In *Aristotelian Society Supplementary Volume,* Vol. 84, No. 2, 2010. Pp. 79-97.

- *Moral Perception*. [Princeton: Princeton University Press, 2013].
- *Moral Perception Defended.* In *Argumenta,* Vol. 1, No. 1, 2015. Pp. 5-28.

Ayer, A. J. *Language, Truth, and Logic.* [London: Gollancz, 1936].

Bakhurst, D. *Particularism and Moral Education.* In *Philosophical Explorations,* Vol. 8, No. 3, 2005. Pp. 265-279.

Berker, S. *Particular Reasons.* In *Ethics,* Vol. 118, No. 1, 2007. Pp. 109-139.

Blackburn, S. *The Oxford Dictionary of Philosophy.* [Oxford: Oxford University Press, 2016].

Brännmark, J, and Lind, A. *Particularism in Question: an Interview with Jonathan Dancy.* In *Theoria,* Vol. 74, No. 1, 2008. Pp. 3-17.

Broad, C. D. *Five Types of Ethical Theory.* [London: Routledge, 1930].

Broome, J. *Reasons.* In *Reasons and Value: Themes From the Moral Philosophy of Joseph Raz.* Ed. Wallace, R. J. Pettit, P. and Smith, M. [New York: Clarendon Press, 2004]. Pp. 28-55.

Cecchini, D. *Problems for Hard Moral Particularism: Can We Really Dismiss General Reasons?* In *Philosophical Inquiries,* Vol. 8, No. 2, 2020. Pp. 31-46.

Copp, D. *Particularism and Antitheory.* In *The Oxford Handbook of Ethical Theory.* [Oxford: Oxford University Press, 2005].

Crisp, R. *Particularising Particularism.* In *Moral Particularism.* [Oxford: Oxford University Press, 2000].

Cullity, G. Particularism and Presumptive Reasons. In *Aristotelian Society* S*upplementary Volume,* Vol. 76, No. 1, 2002. Pp. 169-190.

- *Weighing Reasons.* In *The Oxford Handbook of Reasons and Normativity.* Ed. Star, D. [Oxford: Oxford University Press, 2018]. Pp. 423-442.

Dancy, J. *Ethics Without Principles*. [Oxford: Oxford University Press, 2004b].

- *Moral Perception.* In *Aristotelian Society Supplementary Volume,* Vol. 84, No. 2, 2010. Pp. 99-117.

- *On Moral Properties*. In *Mind*, Vol. 90, No. 359, 1981. Pp. 367-385.

- *Can A Particularist Learn the Difference Between Right and Wrong?* In *The Proceedings of the Twentieth World Congress of Philosophy*, 1999, Pp. 59-72.

- *The Particularists Progress.* In *Challenging Moral Particularism.* Ed. Lance, M. Potrč, M. and Strahovnik, V. [Routledge: London, 2008].

- *Ethical Particularism and Morally Relevant Properties.* In *Mind,* Vol. 92, No. 368, 1983. Pp. 530-547.

- *the Role of Imaginary Cases in Ethics.* In *Pacific Philosophical Quarterly*, vol. 66, no. 1, 1985. Pp. 141-153.

- *Two Conceptions of Moral Realism.* In *Proceedings of the Aristotelian Society Supplementary Volumes,* Vol. 60, No. 1, 1986. Pp. 167-189.

- *Non-Consequentialist Reasons.* In *Philosophical Papers,* Vol. 20, No. 2, 1991. Pp. 87-112.

- *Externalism For Internalists.* In *Philosophical Issues,* Vol. 2, No. 368, 1992. Pp. 93-114.

- *Defending Particularism.* In *Metaphilosophy,* Vol. 30, No. 1, 1999. Pp. 25-32.

- *Caring about Justice.* In *Philosophy,* Vol. 67, No. 262, 1992. Pp. 447-466.

- *Intuitionism in Meta-Epistemology*. In *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition,* Vol. 42, No. 3, 1982. Pp. 395-408.

- *Are There Organic Unities?* In *Ethics,* Vol. 113, No. 3, 2003. Pp. 629-650.

- *Enticing Reasons.* In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* [Oxford: Clarendon, 2004a].

Delapp, K. *Metaethics and the Limits of Philosophy.* In *Filosofiska Notiser,* Vol. 10, No. 1, 2023. Pp. 3-22.

Darwall, S. *Morality, Authority, and Law.* [Oxford: Oxford University Press, 2013].

Enoch, D. *Taking Morality Seriously.* [Oxford: Oxford University Press, 2011].

Finlay, S. *Defining Normativity.* In *Dimensions of Normativity, New Essays on Metaethics and Jurisprudence.* [Oxford: Oxford University Press, 2019].

- *Four Faces of Moral Realism.* In *Philosophy Compass,* Vol. 2, No. 6, 2007. Pp 820-849.

Friedman, J. *Question-Directed Attitudes.* In *Philosophical Perspectives,* Vol. 27, No. 1, 2013. Pp. 145-174.

Gleeson, A. *Moral Particularism Reconfigured.* In *Philosophical Investigations,* Vol. 30, No. 4, 2007. Pp. 363-380.

Hare, R. M. *The Language of Morals.* [Oxford: Oxford University Press, 1952].

Harman, G. *Moral Particularism and Transduction.* In *Philosophical Issues,* Vol. 15, No. 1, 2005. Pp. 44-55.

- *Moral Relativism Defended.* In *Explaining Value: And Other Essays in Moral Philosophy.* [Oxford: Clarendon, 1975].

Hicks, A. *Particularism Doesn't Flatten.* In *Journal of Moral Philosophy,* Vol. 13, No. 3, 2014. Pp. 339-362.

Holton, R. *Principles and Particularisms.* In *Aristotelian Society Supplementary Volume,* Vol. 76, No. 1, 2002. Pp. 191-209.

Hooker, B. *Moral Particularism: Wrong and Bad.* In *Moral Particularism.* [Oxford, Oxford University Press, 2000].

Moral Particularism and the Real World. In *Challenging Moral Particularism.* Ed. Lance, M. Potrč, M. and Strahovnik, V. [London: Routledge, 2008].

Horty, J. *Reasons as Defaults.* In *Philosophers Imprint,* Vol. 7, No. 3, 2007.

Huemer, M. *Ethical Intuitionism.* [London: Palgrave Macmillan, 2005].

Jackson, F. Pettit, P. and Smith, M. *Ethical Particularism and Patterns.* In *Moral Particularism.* [Oxford, Oxford University Press, 2000].

Jackson, N. *Habits and Mental Perspectives: Educating Moral Particularism.* In *The Pluralist,* Vol. 12, No. 2, 2017. Pp. 27-56.

Kagan, S. *The Additive Fallacy.* In *Ethics,* Vol. 99, No. 1, 1988. Pp. 5-31.

Kirchin, S. *Particularism and Default Valency.* In *Journal of Moral Philosophy,* Vol. 4, No. 1, 2007. Pp. 16-32.

Kremm, D, and Schafer, K. *Metaethical Quietism.* In *The Routledge Handbook of Metaethics*. [London, Routledge, 2018].

Lance, M. and Little, M. O. *Defeasibility and the Normative Grasp of Context.* In *Erkenntnis,* Vol. 61, No. 2, 2004. Pp. 435-455.

Leibowitz, U. *Moral Advice and Moral Theory.* In *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition,* Vol. 146, No. 3, 2009. Pp. 349-359.

- *Particularism in Aristotle's Nicomachean Ethics.* In *Journal of Moral Philosophy,* Vol. 10, No. 2, 2013. Pp. 121-147.

- *Explaining Moral Knowledge.* In *Journal of Moral Philosophy,* Vol. 11, No. 1, 2014. Pp. 35-56.

Lerner, A. and Leslie, S. J. *Generics, Generalism, and Reflective Equilibrium: Implications for Moral Theorizing From the Study of Language.* In *Philosophical Perspectives,* Vol. 27, No. 1, 2013. Pp. 366-403.

Lillehammer, H. *Companions in Guilt Arguments in Metaethics.* [London: Routledge, 2019].

Little, M. O. *On Knowing the 'Why': Particularism and Moral Theory.* In *The Hastings Centre Report,* Vol. 31, No. 4, 2001a. Pp. 32-40.

- *Wittgensteinian Lessons on Moral Particularism.* In *Slow Cures and Bad Philosophers: Essays on Wittgenstein, Medicine, and Bioethics.* [Durham: Duke University Press, 2001b]
- *Particularism and Antitheory.* In *The Oxford Handbook of Ethical Theory.* [Oxford: Oxford University Press, 2007].
- *Moral Generalities Revisited.* In *Moral Particularism.* [Oxford: Oxford University Press, 2000].

Lyons, 1977. *Semantics.* [Cambridge: Cambridge University Press, 1977].

Mackie, J. L. *Ethics: Inventing Right and Wrong.* [New York: Penguin, 1977].

MacIntyre, A. *After Virtue.* [Notre Dame: University of Notre Dame Press, 1981].

McKeever, S. and Ridge, R. *The Many Moral Particularisms.* In *Canadian Journal of Philosophy,* Vol. 35, No. 1, 2005. Pp. 83-106.

- *Principled Ethics: Generalism as a Regulative Ideal.* [Oxford: Oxford University Press, 2006a].
- *Turning on Default Reasons.* In *Journal of Moral Philosophy,* Vol. 4, No. 1, 2007. Pp. 55-76.
- *Particularism and the Contingent A Priori.* In *Acta Analytica,* Vol. 21, No. 1, 2006b. Pp. 3-11.

McNaughton, D, and Rawling, P. *Unprincipled Ethics.* In *Moral Particularism.* [Oxford, Oxford University Press, 2000].

- *Contours of the Practical Landscape.* In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy.* [Oxford: Oxford University Press, 2013].

McPherson, T. *What is at Stake in Debates among Normative Realists?* In *Nous,* Vol. 49, No. 1, 2015. Pp 123-146.

Mintoff, J. *Why Moral Principles?* In *Mind,* Vol. 125, No. 500, 2016. Pp. 1133-1159.

Mišcevic, N. *When the Plot thickens: Dancy on Thick Concepts.* In *Challenging Moral Particularism.* Ed. Lance, M. Potrč, M. and Strahovnik, V. [London: Routledge, 2008].

Moore, G. E. *Some Main Problems of Philosophy.* [London: Allen and Unwin, 1953].

- *Principia Ethica.* [Cambridge: Cambridge University Press, 1903].
- *Ethics.* [London: Williams and Norgate, 1912].

Nica, D. *Narrative and Justification in Moral Particularism.* In *Romanian Journal of Analytic Philosophy,* Vol. 8, No. 2, 2013. Pp. 22-32.

O'Neill, Onora. *Towards Justice and Virtue: A Constructive Account of Practical Reasoning.* [Cambridge: Cambridge University Press, 1996].

Price, A. W. *A Quietist Particularism.* In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy.* [Oxford: Oxford University Press, 2013].

- *Was Aristotle a Particularist?* In *Proceedings of the Boston Area Colloquium of Ancient Philosophy,* Vol. 21, No. 1, 2005. Pp. 191-212.

Reiter, R. *A Logic for Default Reasoning.* In *Artificial Intelligence,* Vol. 13, No. 1, 1980. Pp. 81-132.

Richards, R. J. *Evolutionary Ethics: A Theory of Moral Realism.* In *The Cambridge Handbook of Evolutionary Ethics.* Ed. Richards, R. J. and Ruse, M. [Cambridge: Cambridge University Press, 2017]. Pp. 143-157.

Roberts, D. *Why Believe in Normative Supervenience?* In *Oxford Studies in Metaethics,* Vol. 13, 2018. Pp. 1-24.

Robinson, L. *Moral Principles as Moral Dispositions.* In *Philosophical Studies: an International Journal for Philosophy in the Analytic Tradition,* Vol. 156, No. 2, 2011. Pp. 289-309.

Ross, W. D. *The Right and the Good.* [Oxford: Oxford University Press, 1930]

Sandis, C. *Virtue Ethics and Particularism.* In *Aristotelian Society Supplementary Volume,* Vol. 95. No. 1, 2021. Pp. 205-232.

Salay, N. *Thinking Without Global Generalisations: A Cognitive Defence of Moral Particularism.* In *Inquiry: An Interdisciplinary Journal of Philosophy,* vol. 51, no. 4, 2008. Pp. 390-411.

Shafer-Landau, R. *Moral Rules.* In *Ethics,* Vol. 107, No. 4, 1997. Pp. 584-611.

- *Moral Realism: A Defence.* [Oxford: Oxford University Press, 2003].

Shannon, C. *Programming a Computer for Playing Chess.* In *Philosophical Magazine,* Vol. 41, No. 314, 1950. Pp. 1-18.

Sinnott-Armstrong, W. *Some Varieties of Particularism.* In *Metaphilosophy,* Vol. 30, No. 2, 1999. Pp. 1-12.

Singer, P. *Famine, Affluence, and Morality.* In *Philosophy & Public Affairs,* Vol. 1, No. 3, 1972. Pp. 229-243.

Smith, B. *Particularism, Perception and Judgement.* In *Acta Analytica,* Vol 21. No. 2, 2006. Pp. 12-29.

- *Possible Worlds and Moral Philosophy*. In *Teorema: Revista Internacional de Filosofia*, Vol. 20, No. 3, 2001. Pp. 41-50.
- *Particularism and the Space of Moral Reasons.* [London: Palgrave Macmillan, 2011).

Smith, J. *Do Particular Moral Judgements Follow a Rule?* In *The Southern Journal of Philosophy,* Vol. 40, No. 2, 2002. Pp. 269-294.

Sneddon, A. *Thick Concepts and Holism about Reasons.* In *The Journal of Value Enquiry,* Vol. 44, No. 4, 2010. Pp. 461-468.

Stangl, R. *Particularism and the Point of Moral Principles.* In *Ethical Theory and Moral Practice,* Vol. 9, No. 2, 2006. Pp. 201-229.

- *A Dilemma for Particularist Virtue Ethics.* In *The Philosophical Quarterly,* Vol. 58, No. 233, 2008. Pp. 665-678.

Strahovnik, V. *Defeasibility, Norms and Exceptions: Normalcy Model.* In *Revus – Journal for Constitutional Theory and Philosophy of Law*, Vol. 29, No. 1, 2016. Pp. 61-76.

Thomas, A. *Another Particularism: Reasons, Status and Defaults.* In *Ethical Theory and Moral Practice,* Vol. 14, No. 2, 2009. Pp. 151-167.

- *Practical Reasoning and Normative Relevance: A Reply to McKeever and Ridge.* In *Journal of Moral Philosophy,* Vol. 4, No. 1, 2007. Pp. 77-84.

Väyrynen, P. *Particularism and Default Reasons.* In *Ethical Theory and Moral Practice,* Vol. 7, No. 1, 2004. Pp. 53-79.

- *Moral Generalism: Enjoy in Moderation.* In *Ethics,* Vol. 116, No. 4, 2006. Pp. 707-741.
- *Grounding and Normative Explanation.* In *Aristotelian Society Supplementary Volumes,* Vol. 87, No. 1, 2013. Pp. 155-178.

Vervidakis, S. *Varieties of Quietism.* In *Philosophical Inquiry,* Vol. 30, No. 1, 2008. Pp. 158-175.

Westing, A. H. *A Note on How Many Humans That Have Ever Lived.* In *BioScience,* Vol. 31, No. 7, 1981. Pp. 523-524.

Whiting, D. *Defending Semantic Generalism.* In *Analysis,* Vol. 67, No. 04. 2007. Pp. 303-311.

Wittgenstein, L. *Tractatus Logico-Philosophicus.* [London: Routledge, 1922].

- *Philosophical Investigations*. [Oxford: Blackwell, 1953].

Zamzov, J. *Rules and Principles in Moral Decision Making: An Empirical Objection to Moral Particularism.* In *Ethical Theory and Moral Practice,* Vol. 18, No. 1, 2015. Pp. 123-134.

Zingano, M. *Moral Particularism and Aristotelian Ethics.* In *Nature and the Best Life: Exploring the Natural Bases of Practical Normativity in Ancient Philosophy.* Ed. Rossi, G. [Hildesheim, Zürich, New York: Georg Olms Verlag, 2013]. Pp. 95-126.