

Value and Variance: Dancy, Moore, and Scanlon.

PETER NICHOLAS TUCK

2021

Value and Variance: Dancy, Moore, and Scanlon.

PETER NICHOLAS TUCK

**A thesis submitted in fulfilment of
the requirements of Manchester
Metropolitan University for the
degree of Master of Arts by Research**

**Department of Philosophy
Manchester Metropolitan University
2021**

Contents

Title	1
Contents	3
Abstract.....	4
Introduction	5
Chapter One : Non-Natural Realism	10
Chapter Two : Moore and Scanlon.....	32
Chapter Three : Dancy and Scanlon	53
Chapter Four : Intrinsic and Default Values	68
Chapter Five : Organic Unity and Value Holism	80
Conclusion	88
Bibliography	91

Abstract

This thesis deals with two important but often conflicting strands in ethical theorising. The first of these is the analysis of what the terms 'good' or 'valuable' really refer to at all, while the second deals with how what is really good or valuable can change, depending on context. While analyses of value often privilege one strand at the expense of the other, this thesis considers two popular approaches to how the second sort of view can in fact be reconciled with the first: Jonathan Dancy's Value Holism, and George Moore's Organic Unity.

While these two approaches, which deal with different parts of the real-yet-contextual problem, are often held to be incompatible, this thesis argues that accepting a particular view about the first question, namely Thomas Scanlon's buck-passing view, allows us to go on to accept both Dancy and Moore's views about the second.

Since Dancy's view contextualises the real goodness or badness of individual things, while Moore's view contextualises the value of groups of valuable things, or even good and bad things together, the novel argument in this thesis that allows the acceptance of a combined view promises to provide a complete and internally consistent 'holistic' analysis of value.

Introduction

There are three common and apparently broadly shared assumptions that we might describe as both commonsensical, and pre-theoretical. The first of these assumptions is that some form of what philosophers would call 'evaluative realism' is true: that there are objective truths about what is good and bad (at least in some cases), and that, because of this, we are capable (at least in some cases) of knowing what they are. In this sense we might hold that things like justice, truth, freedom, or equality are *actually good*, outside of any particular personal viewpoint, place, or time. It is because this sort of idea is so widely found throughout different places and times, and so commonly leveraged or applied in ordinary discussion, in society, and in politics, that it is a fitting informal starting point for philosophers to seek to explain formally.

The second of these assumptions is the one that philosophers like to talk about using terms like 'intrinsic', 'instrumental', 'final' or 'conditioned' value: that if evaluative realism is true, it will also be true that good things are good in different ways, and/or in different amounts. These are technical terms with various relevant definitions. For example, intrinsic value is meant to be the value something has because of its own attributes alone. Instrumental value is where something is valuable because of the use it can be put to. Final value is meant to mean that something is a fitting end to

aim towards, rather than merely a way to achieve that end. Finally, conditioned value is where the value of something is conditional based on the value of another thing. This is predicated on the first assumption discussed in the previous paragraph. An obvious example to explain this point is to imagine asking someone if a baby or a pencil are good respectively, and then asking if they are both as good as each other. We would expect to hear not only that the baby is much *more* valuable than the pencil (amount) but that the baby actually has a different *type* of value than the pencil, one that seems fixed and immutable in a special way. We might say that it is impossible for the baby to be anything other than valuable, while we wouldn't say the same thing for the pencil. The point here is that there is common-sense motivation for discussions surrounding intrinsic vs instrumental values, means vs ends, and so on: again, the task for the philosopher is to make coherent sense of this in theory.

The third and final of these assumptions is that there is something *contextual* about value. In this sense, even if value realism is true, it will still be the case that differences in places, times, persons and other contextual factors may influence what is actually good or bad *here*. However, when considered theoretically and in line with what various thinkers have said on the subject, for this to be non-trivial it must be possible for value that is usually *fixed* to be able to change, whether in amount or actually type of value.

This point is predicated on both the first and second assumptions I pointed out: we need to think there are some facts about which things are valuable, in what way, and in what amount, to think that there is any fuss at all about claiming that these things can change. If we were to hold, for example, that facts about what is valuable

are really just facts about what things some particular person believes, or feels, are valuable, then we could happily accept that these could change, and there would be no case to answer.

Similarly, even when we think there are such facts, we seem to think they are limited: we might think there are facts about equality, for example, being valuable in a special way, but not think the same for paperclips at all. Therefore again, for the idea that value is contextual to actually be interesting, non-trivial and non-obvious, we need to be able to say how we think something that really is definitely valuable can be contextualised.

Therefore, while this thesis will consider recent philosophical approaches that proceed from all three of these starting points, it will necessarily do so in order. Regarding our first question, I must distinguish a particular point of order from the outset. The kind of metaphysical view implied by that question obviously points towards a particular philosophical term: realism, more specifically, moral or evaluative realism. However, this is an umbrella term under which a variety of distinct views are corralled. While there are various ways they can be taxonomized, the simplest and most relevant distinction for our purposes is that between *natural* and *non-natural* accounts. The second of these will be the topic of this essay.

So, I distinguish the question 'is non-natural realism true?' from the question 'If non-natural realism is true, how well have philosophers done at explaining it?'. For this reason, the thesis will not directly survey opposing views that hold that value is

entirely natural, or entirely unreal. These include the full range of reductive realist, non-cognitivist, and anti-realist positions.

The first chapter will consider what standards we ought to have for non-natural realism, primarily via the extremely relevant criticisms of J.L. Mackie, who has given arguments that problematise the very existence of evaluative facts; or at least the explanations philosophers have given in support of them. There are two ways Mackie's views can be taken. The stronger way is to conclude that he actually refutes the possibility of non-natural realism, or indeed any shade of realism. The weaker way is to conclude that there is still hope for realism, but that Mackie's conclusions have to be taken seriously in terms of setting limits on what sensible approaches might look like. This essay will argue on this second sort of line, and explain how Mackie demonstrates that all realist views must meet certain standards including clarity, explanatory potential and relative parsimony if we wish to even consider endorsing them.

The second chapter will examine G.E Moore's view on our first question: what it means to say something is valuable at all. In doing so, it will consider criticisms of Moore and how his view might be improved. Writing at the turn of the 20th century, Moore was the first in the modern Anglo-analytic tradition to consider what philosophers should say about our three questions, and his work sets the stage for later contributions in the same line. By considering these criticisms, the chapter will come to introduce and explain Thomas Scanlon's answer to the same question, which comes much more recently and incorporates those problems, and conclude that it is superior to Moore's on certain central considerations.

The third chapter will examine a third view: that of Jonathan Dancy. Dancy positions himself between Moore and Scanlon, and while he agrees that Scanlon's view is better than Moore's, he claims his own revisions are required due to two key problems. I argue that both of these concerns are misplaced, and therefore Scanlon's view is ultimately the most attractive. At this point, I will have moved from a general assumption that evaluative realism might be true to endorsing a specific, considered formulation of the point. By doing this, I am able to move to consider our second assumption in a similarly precise and considered way.

The fourth chapter therefore considers what philosophers have to say about our second question: whether there are two or more sorts of value. I take an ontological, grounded approach, focusing on an intrinsic/extrinsic distinction that sorts valuable objects based on whether they are such because of their own features. I then introduce a second, similar notion: that of default or non-default values, or reasons. I compare the two notions, and explain how the second sort, that of a default, is more directly applicable to the comparison I will make in chapter five. Again, by carrying out this interrogation I am able to move from a general assumption about sorts of value to endorsing a specific theoretical way of explaining the truth of such. This allows me to move to consider our third assumption: that such default values can nevertheless change in various ways.

In the fifth chapter, I directly compare two philosophical views on that our third question: whether valuable things *of the same sort* can change in status, amount, or valence. The first view is that of Moore, who explains how multiple objects that are good or bad can interact to create such contextuality. Jonathan Dancy takes the

opposite approach: his view predominantly tells us how *a single thing* can change. It seems promising that both views, when taken together, might provide a more complete explanatory schema than either alone. However, several contributors to the subsequent literature point out two fundamental places the two views seem to be incompatible. I argue that these apparent conflicts can be overcome, and specifically, that my earlier endorsement of Scanlon's view about value gives us extra resources to do so.

This is the specific and narrow claim this thesis is oriented towards that is an original contribution to knowledge: Not only is it possible to accept both Moore and Dancy's views about the contextuality of value, but it is accepting Scanlon's view about what value is in the first place that allows us to confidently do so, to our benefit, as it offers the most complete and attractive approach currently available in the literature towards theoretical explanation of our pretheoretical assumptions.

Chapter One : Non-Natural Realism

In the introduction I identified three plausible pre-theoretical intuitions about axiology as promising starting points for theoretical discussion, and promised to take them in order. The first of these was that some form of 'non-natural' realism is *prima facie* worth investigating. Before moving to examine specific realist theories, I will start by defending a set of minimum standards that any such theory ought to meet: I will then be in a position to compare some of the most promising theories according to these standards.

I will do this via the vehicle of examining relevant criticisms of realism: objections that realists ought to take seriously, and ones that limit and guide the range of possible ways of explaining the metaphysics of non-natural realism. Before I move to these criticisms, however, I will have to more clearly define several key terms that will be relied upon: non-natural realism, normativity, naturalness, and grounding.

Firstly, what consensus is there on the basic commitments of non-natural realism itself, in the theoretical and evaluative sense? Philosophers have broadly coalesced around two key theses (Finlay, 2007):

1. There are objective facts about what is good or bad, outside of any particular place, time or viewpoint. In this sense, evaluative truths are *mind independent*. This is a general feature of any realist view e.g. mathematical realism would hold that a triangle having three sides is a fact that would hold outside of any place, time, instantiation or viewpoint: the term 'triangle' refers to a particular abstract but real thing. This is the *metaphysical* thesis. For brevity, from now onwards I will generally use the straightforward term 'robust' to refer to this putative existential status of non-natural facts or properties.
2. Following from 1, evaluative claims will always be true or false (*truth-apt*) and some of them will be true. For example, for any particular *x*, if person *a* says *x* is good, and person *b* says *x* is bad, and *x* is in fact good, then person *a* will have spoken truly, and *b* falsely, with this being guaranteed by 1 in a binary sense. This is the *epistemic* thesis.

It seems reasonable to say here that on this view, the epistemic thesis is partially guaranteed by the metaphysical one, but not the other way around. Once

one accepts the plausibility of the first sort of claim, one is still left with the question of exactly how an agent forms the correct justified belief in relation to the corresponding evaluative reality, and there are multiple strands in the literature here too. However, a survey of the debate on realist epistemology still leaves an obvious further question: what is the shape of this thing that we can know? What is it about *the way the world is* that lies behind, and supports, our correct epistemic practice?

The metaphysical claim 'underwrites' the epistemic one: accepting it means accepting that there are true or false ethical or evaluative claims and accepting at least the *possibility* of a correct account of belief formation: but this conclusion isn't exclusive to those who hold 1. What is unique to robust realism is its ontology: one can hold that evaluative statements can be true or false while denying that their truth-makers are robust, distinct, normative kinds. I want to interrogate what robust realists say about the nature of these truth-makers in particular.

To do this, I will have to explain something the separate definitions of, and relationship between, the "normative" and the "natural", the debate mainly turning on exactly how the relation is analysed. I will have to do so because the main criticisms of realism, as well as the range of views on offer, make much of this distinction.

Normativity, as the name suggests, refers to norms of, for example, action, attitude, or thought. As an umbrella term normativity can include legal norms or social norms, but I distinguish this broader use of the term which includes things that we usually think are socially constructed from a narrower use that accords with the definition of realism above. While we might argue that, for example, traffic regulations or social norms about polite behaviour can be justified by reference to more

fundamental evaluative or moral facts, it is those more fundamental facts that we are actually claiming are independently robust: the broader sort are simply applications of this narrower core. It is evaluative or moral norms we have to directly justify: the strength of the justifications we can give for social or legal norms are indirectly based in the strength of those fundamental justifications.

On this narrower, robust and ethical definition normativity is often described in terms of 'should' or 'ought' in the prescriptive or *binding* sense, but to do so is to, in my view, inappropriately suggest that discussion of the normative is discussion of the *deontic*, which is a subset of normativity and by definition surveys that which is binding¹. Deontic concepts include rightness and duty. The normative realm is usually understood to include other, broader families of concepts than the deontic, most notably evaluative (including ethical and aesthetic), but often also of reasons, of fittingness, or even of logical² or epistemic³ content.

Moreover, there are various views that hold that the normative 'should' is sometimes only in the 'recommending' or 'enticing' business (Dancy, 2004a, 2004b, Little, 2013): e.g., there are many things that make liking the Mediterranean enticing, and many that recommend visiting it: neither sort are binding. For this reason, I am inclined to work with the definition of normativity as the study of how we *should* (whether in the prescribing or enticing sense) act, think, or feel, whether binding or

¹ Ancient Greek, *deon*, binding.

² For example, that logic *prescribes* for us, or *entices* us, to reason or infer in a certain way.

³ In the sense that epistemological realism involves defending norms about correct and incorrect belief formation.

not, or indeed whether 'correct' or not. While there is indubitably a normative 'should', it seems reasonable to allow that there can also be a normative 'could'.

As referred to in the definition of realism above, robust realists are committed to a conception of *normative* truths as mind-independent and not spatiotemporally located. As well as being abstract (Mintz-Woo, 2017) they are held to not possess causal powers (Bedke, 2009): they are causally 'impotent'.

Debbie Roberts (2018, 4-7) provides a partial survey of what is meant philosophically by the 'natural'. There are some obvious tokens of the type that are uncontroversially natural (e.g., if anything exists naturally, dogs, trees and stones do) but for our purposes we want a hedged description of the type itself, and specifically one that clearly demonstrates a difference with normative content, for reasons that will become apparent.

An ontological approach neatly identifies two characteristics of the natural that are directly opposite to those two characteristics the normative is held to have: that it is spatiotemporally located or concrete, and therefore that it has causal powers (Bedke, 2009, 189).

A Semantic approach aims to parse the two sorts of things by pointing out that we make descriptive statements about the natural, which contrast to the prescriptive ones we make about normativity. For example, if I point at a dog and say 'this is a dog' I am giving a completely uncontroversial description of the thing. If I then say 'this dog has moral value' I am making a distinctly prescriptive, normative claim (Roberts, 2018, 4-7).

Epistemological approaches categorise the natural as that which can be studied using some particular method of investigation, for example that naturalness is that which can be studied by the sciences, or the empirical method. While originally philosophers preferred to hedge this as the *natural* sciences (e.g., Moore, 1903) recently there has been a move to expand the definition to include the social sciences (e.g., Shafer-Landau, 2003, 19).

Naturally, all of these ways of distinguishing the natural are disputed (McPherson, 2015). However, my primary objective is to distinguish the natural from the non-natural by way of emphasising the important familiarity of the natural. For our purposes, a packaged combination of the above views seems to have this virtue. By pointing out that natural kinds are concrete and spatiotemporally located it emphasises a familiar mode of existence. By including a semantic distinction, it distinguishes the way we consistently describe the two kinds. By including an epistemological clause, it points to a familiar and widely trusted way of accessing facts about the sort of things identified by the ontological view: empirical investigation.

Naturalness, then, can be understood as referring to things that exist in spacetime, have causal powers, can be defined using descriptive language, and can be apprehended by the epistemic practices of the contemporary branches of science. I feel that this definition is relatively theory neutral and is accommodating enough that it will make room for the different ways the philosophers I discuss invoke the term.

So, we have two sorts of facts about real and familiar objects. Again: if I say that dogs are good, that they have value and worth, and that this has something to do with what a dog is in the first place, I am referring to two things: all the *natural*

facts that make the dog what he is, and the further, apparently *derived, normative* fact: that the thing that possesses those natural properties is good.

It amounts to orthodoxy in ethics, or normative philosophy more generally, that there must be some consistent relationship between the normative and the natural (Roberts, 2018). This can be motivated by appealing to the view that if there is a normative truth about any natural thing, e.g., that a particular object is good, then any truly identical object will also be good. This fundamental point was originally made in the recent literature by R.M. Hare (1964, 80-81), where he writes:

“Suppose that a picture [P] is hanging upon the wall and we are discussing whether it is a good picture... Suppose that there is another picture next to P ... (I will call it Q). Suppose that either P is a replica of Q or Q of P, and we do not know which... Now there is one thing that we cannot say; we cannot say ‘P is exactly like Q in all respects save this one, that P is a good picture and Q not’.

If we were to say this, we should invite the comment, ‘But how can one be good and the other not, if they are exactly alike? There must be some further difference between them to make one good and the other not.’ Unless we at least admit the relevance of the question ‘What makes one good and the other not?’ we are bound to puzzle our hearers; they will think that something has gone wrong with our use of the word ‘good’.”

This last remark about the reaction we can expect from our (presumably reasonably intelligent, rational, and thoughtful) interlocutors, taken together with the one about the limits on our own speech, are meant to highlight the putative status of this relation as a *conceptual* truth. A conceptual truth is meant to be a concluding

truth, predicated upon propositions about certain concepts, where correctly grasping and understanding the relevant concepts present in the argument results in one fundamentally appreciating the resulting conclusion. While this example is analogous in that it discusses aesthetic value, consider the following pairs of statements about specifically evaluative and moral claims:

1. Dogs are good but cats are bad.
2. It is wrong to steal from small family-owned businesses, but permissible to steal from large corporations.
1. Dogs are good, but dogs are bad.
2. It is wrong to steal from small family-owned businesses, but permissible to steal from small family-owned businesses.

There is something fundamentally puzzling about the second pair in that they posit a normative difference without a natural difference. It seems axiomatic to say that when discussing value in the realist sense, all the things *in the world* we want to be able to say are good or bad (dogs, cars, political parties, drugs) will be good in a way that intelligibly relates to the way they otherwise are: this will be the reason why goodness, defined earlier as an abstractum, outside of time and space, attaches itself to particular concreta (Mintz-Woo, 2017, 707) in particular places and times. There will be some natural features, or collections of such, that goodness or badness relate to case by case. This point can be simply formalised as 'no change in B without change in A', where A refers to some natural fact(s) and B refers to some normative fact(s).

Now there are two competing approaches in the 'robust' literature regarding how precisely to further develop this consideration. The first way philosophers like to

describe this relation is as *grounding*. Grounding (Berker, 2018), as the name implies, refers to truths about what is good being grounded in truths about *the way the natural world is*, and has also been referred to as emergence, or resultance, (Dancy, 1981, Potrc, 2004) in that these truths *emerge* or *result* from the way the natural world is. The simplest way to explain what is meant here is to follow Selim Berker in pointing out that the relation is accurately captured in the ordinary English phrases *in virtue of*, or *because of* (2018, 730).

To again return to our dog example: it will be a collection of natural facts e.g., that he is a mammal with a functioning brain and nervous system, that he has certain perceptions, certain experiences, certain activities he carries out, and so on, that all together *ground* the normative fact that he has value. To give a supporting analogy about action: if the dog is clearly thirsty and I discern that I *ought to* give him water, this invisible normative *ought* will arguably be instantiated *because of* facts like the dog's biology requiring him to drink water to survive and thrive, the current amount of water in his body, the fact that he is my dog and I am the only one around to give him water, the fact that he does not know how to use taps, etc.

The grounding view seems to follow from Hare's compelling observation: if we concede that for something to become less valuable it must change naturally, we also concede that a natural change can instantiate an increase in value. In Hare's example, presumably vandalising one of the identical paintings would support its decreasing in value, whereas restoring it physically would also restore its value. If *adding* something physically can cause something to be added normatively, then natural features must

really have good-making or bad-making qualities: it is good or bad *in virtue* of those qualities.

The second way philosophers like to characterise this relation is as supervenience (Kim, 1990, Roberts, 2018), which refers to a combination of two views, the first of which is simply the grounding view described above. The second, unique part that is added on (Berker, 2018, 733-737), is the doctrine of necessary bilateral covariance of the form:

No change in A without change in B, no change in B without change in A.

The problem with supervenience is that while the first half, the 'grounding' half of the view, is comparatively uncontroversial and theory neutral, the second part, the 'necessary covariance' part, implies certain assumptions that directly contradict opposing views I will examine regarding my initial third point: contextuality (Roberts, 2018, 6). In this sense, it presupposes answers to certain questions *beyond* the one at hand, the one about the relation between the normative and natural, and loads them into the initial view by definition.

Necessary covariance assumes the truth of an approach to ethics that, while dominant throughout the philosophical tradition, is challenged by just the sort of newly persuasive views about contextuality I will go on to discuss. It only allows what are called, respectively, 'atomism' and 'generalism' (Dancy, 2000) about value. This kind of view holds that normative truths hold invariantly across contexts, specifically in that for any particular arrangement of natural facts, there is a fixed corresponding normative truth (Roberts, 2018, 6).

'Holism' and 'particularism' (Dancy, 2004b), the opposing counterparts of atomism and generalism, are parts of a view which holds that secondary 'enabling' factors (like facts about people, places or contexts) *that are not themselves the 'grounds' for those truths* can strengthen, weaken, or reverse the polarity of normative truths, even when all the same central natural facts are present: to select any hedged batch of natural facts, and to pair them with a particular normative fact (perhaps packaged and codified as a normative principle) will be to ignore a potentially persuasive position which holds that there is no limit to what kinds of natural facts can weigh on any particular normative truth: grounding is always *potentially* global (Dancy, 2003, 2008), rather than local in a clearly demarcated way.

While I will return to these claims in fine detail during my later discussion of contextuality, suffice to say at this point that it will be better to discuss 'grounding' while leaving aside the specific supervenience requirement of necessary covariance. It is necessary to make this distinction at the outset because in the recent 'metanormative' literature the supervenience term is often used in place of the grounding term when entailment is really the relevant topic at hand, not covariance (Berker, 2018). Therefore, at any point where I discuss the work of thinkers who discuss 'supervenience', I should be read as discussing their claims insofar as they refer the grounding relation, treating the necessary covariance bit as orthogonal.

With these terms in hand, I now turn to the main point: strong criticism of robust realism. Perhaps the best known, and certainly *prima facie* persuasive, criticism of realist metaphysics is J.L. Mackie's 'queerness' objection (1977, 38-42). The claim, simply put, is that realism invokes a certain amount of metaphysically mysterious 'stuff'

that is objectionably different to everything we otherwise know to exist: specifically, the natural world that (as above) we have strong, uncontroversial, and independent knowledge of, that the normative supposedly emerges from.

Now, while the queerness objection encompasses several different worries about realism that can be unpacked, a distinct claim that it makes explicit from the outset is that realist metaphysics require that the 'queer' thing, additional to any other way of interpreting 'queerness', is queer in that these properties are meant to *prescribe* to us, when nothing 'natural' in the universe seems to do so.

This seems to be a variation on the traditional Humean challenge, where doubt is cast upon moving from a descriptive *is* to prescriptive *ought*, and the complaint here is that the realist basically conjures the ought into being, building it into an invisible 'normativity' by definition, while claiming it is entirely *grounded* in, and therefore predicated upon, that natural world. This seems to be supported by a distinction realists are happy to allow: that natural objects that exist concretely in spacetime have causal powers in the usual way, while normative abstracta ought not to. Realists have to explain how goodness, for example, can *cause* you to respect or appreciate it where it is instantiated, or further, to go on to act in a certain way in light of such respect or appreciation. There is, however, an obvious rejoinder, which aims to show that this is no special worry for robust realism, which the queerness objection is meant to be specifically problematic for.

The first point to make is that the criticism holds just as well for naturalists too. While naturalists don't posit any of the mysterious metaphysical stuff that inspires the queerness objection, they simply displace the *prescriptive* dimension of normativity

onto natural/descriptive facts or properties. For the reductive naturalist who is still a realist, who thinks there are independent moral facts, it must be a natural fact or property, or collection of natural facts or properties, which prescribe to us: this is arguably just as mysterious. The real problem, then, is that morality *prescribes* to us at all, whatever theory we endorse.

This part of the objection has an initial 'knock down' flavour: if it holds, then it holds for all kinds of normative realism and not just value: it would be superfluous to consider what the objection more broadly has to say about the specifically evaluative metaphysics realists advocate for. However, the second point I will make is one that defends all kinds of normative realism (Bedke, 2010) from the prescriptive part of the objection, and allows us to consider the rest of it on its own terms. This is the view that normativity is *rational*, and that it is not normativity that prescribes to us, but *rationality* that prescribes to us, and a failure to accept the prescriptive part of evaluation, as far as it goes, is no different to failing to accept any other sort of rational normative conclusion, whether evaluative, logical or epistemic.

For example, consider again mathematical realism. The fact that a triangle has 3 sides could be said to be prescriptive, in that it seems to demand a particular response: that you believe triangles have 3 sides. Someone who, upon considering triangles, thought instead that they have 2 or 4 sides, would ordinarily be described as 'irrational': she has failed to meet a rational demand. Moore, drawing such an analogy between moral and mathematical realism, observes that someone who does not accept triangles have 3 sides does not understand what it is to be a triangle (1903,

11). The point here is that mathematical truth, in the realist sense, always rationally requires, or *prescribes*, for us to accept that truth.

To see this point about truth made more directly, consider epistemic realism itself. For the epistemic realist, presumably *truth* itself is prescriptive, according to whatever specific, putative, epistemic norms that realist advocates for: someone who violates epistemic norms could fail to reach the correct conclusion, which, when apprehended, is found to be compelling. This means that if someone was adamant in holding that the earth is flat, that the Queen of the United Kingdom is a lizard, or that the coronavirus is caused by '5G', they would ordinarily again be described as irrational, and perhaps again, to paraphrase Moore, as not understanding what it is to be *true*.

Returning to the topic at hand, evaluation, we can see what kind of position we should want to take. Those who do not recognise the value of a human infant, or the need of a thirsty dog, are simply not recognising that these things call for a certain response in terms of that rational *should*, in virtue of all the facts about them: the baby for an apprehension of value, the dog for water. They are not truly appreciating what it means to be a helpless baby, or a thirsty dog.

It seems reasonable therefore to say that the 'prescription' dimension of Mackie's objection is not clearly a 'knock down' argument, and that it can be addressed by arguments found in the 'companions in guilt' (Brink, 1984, Lillehammer, 2019) family, maintaining that real evaluative and moral truths prescribe to us not just because they are evaluative or moral, but because they are real and true. According to the previous sentence, it seems that this objection begs the question against

realism: reality and truth are baked into the definition itself, the thing that is being argued against, and it seems to presuppose that the thing at hand does not exist in the way it is defined. Insofar as it argues that moral facts do not possess features that all other sorts of facts do, Mackie's argument really seems to imply that they are not facts from the get go, and boils down to a tautologous conclusion of the form 'moral facts cannot be facts because moral facts are not facts', insofar as real facts are compelling.

If we adopt this companions in guilt (Brink, 1984, Lillehammer, 2019) type picture of different categories of normativity being relevantly congruent in their rational reality, then questions about someone failing to accept that they *should* really think x is good, or that they should believe x, are arguably questions for evaluative or epistemic *psychology*, respectively: but not directly for metaphysics.

It is important to note that the discussion above is not meant to show that causation is no problem for the robust realist at all: just that it is no unique or special one. However, when I compare specific realist views, it will be possible that some of them have unique problems that weigh on causation and have to be individually dealt with: but the problem will be with the specific view, not with the notion of compellingness itself.

However, there is more to the queerness objection than prescription, or causation. The first problem arises from the fact that irreducibly normative truths are held to be causally inert because they are abstract universals, and whether prescriptive or not, Mackie complains, seemingly correctly, that the metaphysical 'stuff' robust realists advocate for is objectionably strange, and unlike the complaint about

prescription, this *is* one that weighs specifically on the non-naturalist: whether one is a non-naturalist about ethics or mathematics.

The reductive naturalist has the advantage of only relying on familiar natural kinds, while the robust realist defends an entirely separate and unfamiliar category of normative kinds. Companions in guilt strategies are not obviously applicable here also: while I mentioned above that this complaint weighs specifically on non-naturalists generally, in practice it may be easier to accept more familiar abstract logical or mathematical truths over evaluative ones here, and it seems an explanation is called for. So, it seems to me that this is not an objection that touches upon existence directly, but only indirectly. What it is really about is satisfactory explanation of what is being posited: so, the form of the argument is something like:

1. Extant approaches to normative metaphysics do not *explain* their 'metaphysical stuff' clearly enough (according to some desideratum or other).
2. Without a suitably clear enough explanation, there is no strong reason to believe the 'stuff' exists.
3. The stuff does not exist.

My purpose here is simply to establish a particular desideratum for realist theories: that of *clarity*. This objection can only be evaluated further in relation to some specific realist view, and when I later compare several such views, I will evaluate the strategies they respectively use to render their specific formulations clear and unmysterious.

This is the first way in which Occam's razor can be applied to non-naturalist theories: by positing a second *fundamental* normative category, realists are committed

to explaining the existence of what they invoke in line with the strangeness of it, and when taken into account while constructing metanormative views can be called *qualitative* parsimony (Morton & Sampson, 2014, Sendlak, 2018).

It is worth noting that Mackie's frustration is justifiably grounded in the literature insofar as he is responding to non-naturalists (e.g., Moore) who ostensibly argue that non-natural entities are undefinable, irreducible, unanalysable in natural language, etc. and that it really may be the case that realists have failed to deal with their own self-inflicted problem. My task then, will be to demonstrate that undefinable does not mean completely indescribable, and that the philosophers whose views I defend have in fact done enough to clarify their positions.

The third part of the queerness objection, it seems to me, is one of *amount* (Turner, 2016, 380-381). While Mackie's formulation of the point doesn't directly discuss this, it is one that reasonably entails. If whatever 'metaphysical stuff' at hand is unacceptably mysterious, matters will hardly be improved by adding more of it. Reasonably, if there were two realist theories that were both deemed equally mysterious as far as it goes, and one of those theories, *ceteris paribus*, made twice as many independent metaphysical claims as the other, someone otherwise persuaded by the queerness objection ought to think it twice as problematic (Turner, 2016, 380-381). This is the second way Occam's razor can be applied, and when taken seriously while developing metaphysical views can be positively called *quantitative* parsimony (Baker, 2003, Sendlak, 2018).

This consideration is what makes a more than century-old programme in analytic ethics attractive (Wodak, 2019): defining one normative category in terms of

another that is argued to be more *normatively fundamental*. Consider the different variations of the relation between the normative and the natural discussed previously, together with my discussion of queerness. For any realist who isn't a reductive naturalist, that is, one who argues for a complete metaphysical and explanatory *reduction* of the normative to the natural, *at least* one normative concept has to be defended on its own merits. For the reductive naturalist, there seems to be no trouble in how many fundamental normative 'bits and pieces' you have: what will be at stake is instead how well you appeal to their natural home: you can have as much stuff as you can reasonably give a natural explanation for.

However, for the other sorts of realist, we have demonstrated how the queerness objection requires them to be thrifty in how much metaphysical stuff they have that is fundamental. I say this to make it clear than when we talk about 'amount', it is not in the sense of the number of separate concepts or terms that are used, but in the ontological and categorical sense. What this means is, if a successful argument for some conception of goodness is given, hundreds of specific claims can then be made about specific tokens of the goodness type (e.g., charity, fairness, kindness) that are grounded in the original umbrella claim without incurring any extra metaphysical debt (Turner, 2016, 380-381).

This is why this broader part of the objection, unlike the claim about prescription, is specifically problematic for robust realists and one they must account for: as has been made clear, reductive naturalists argue for one, uncontroversial ontological category (the natural) while non-naturalists invoke at least two, the second of which *is* controversial.

As well as being selective in terms of how well they can explain *it*, parsimony entails they must be selective regarding much explanatory power the theory has: what *it* can explain. On this last point, it seems reasonable to say that, at least in terms of 'metaphorical weighting', each extra 'bit' of metaphysics should come with a 'bit' of explanatory power of *at minimum* 'equal' weight.

The required extra explanation that ought to come with extra metaphysics comes in two sorts, one weaker and one stronger. The first one is the *amount* of explanation, as described above in terms of 'metaphysical weighting'. The point here is that theories will be compared like for like *on the same question*, and the winner will be the one that can say more about *that* question. The second point is one about *unique* explanation: where one theory can answer a question that another doesn't seem to be able to answer at all. We can call these superior explanatory power, and unique explanatory power, respectively.

The point here is that a promising strategy for the realist is to pick one normative category to explain the others in terms of, and that the choice of category should be one that the realist can most promisingly both explain and use to explain. Before I say what this might be, I will briefly turn to a slightly different but related topic: what ways might we have of selecting this fundamental normative notion, irrespective of what has been said above? So far, I have briefly explained why such a strategy might be attractive, but said nothing about how to accomplish it past this starting point.

The first way involves the notions of priority, and potentially, of parity. I will start with priority. It immediately seems that there are two related ways one normative

concept can be found to be 'prior' to another: metaphysical priority, and logical, or explanatory, priority. Metaphysical priority refers to the grounding relation discussed above: that one normative notion obtains *in virtue of* or *because of* another. For example, if consequentialism is true, it would seem that the rightness of some particular act obtains *in virtue of* that act having the best consequences.

Another version of the point that holds more fundamentally involves logical, or explanatory, priority. This sort also implies the first if the referents of the premises and conclusions are all things that are independently held to robustly exist: cause and effect implies ontological dependence, and therefore grounding, here.

This means that if an argument applying some normative notion (concluding with a claim about that notion) is predicated on premises that refer to other normative notions, and most obviously at bottom to natural notions, there seems to be a logical entailment. E.g., if there is an argument of the simple form:

If A, then B. If B, then C.

Then if the argument otherwise meets all the usual standards arguments are judged by, the form of the argument will be a point in favour of the real priority of A over B, and B over C. for example:

1. There are certain 'natural' properties that give us reasons to think the thing that has them is 'valuable'.
2. The property of being a conscious being (however we construe or measure this) is such a natural property.
3. We therefore have reasons to think conscious beings are valuable.
4. Unless there are stronger reasons *not* to think they are valuable.

5. There are no stronger reasons against their being valuable.
6. Conscious beings are valuable.

This example argument can obviously be criticised and doubted in all the usual ways any argument fairly can. However, the point is that if the argument is otherwise judged both sound and valid its form will be an *additional* 'point' in favour of the priority, and therefore fundamentality, of the referents of the premises over the conclusions. For example, the argument above, following a form similar to that implied by the discussion so far, takes the normative notion of 'being a reason for' (for whatever reason) as central, claims it has some rational relation to the way the world is (as discussed earlier) and aims to demonstrate success in using it to analyse a second normative notion: here, goodness. It argues from natural properties, to the normative property of a reason, to a conclusion that defines 'valuable' in terms of the first two, in that order.

The third sort of view I mentioned is one that includes parity. In this sense, one might argue that two or more normative notions (for example, goodness and rightness) are utterly different, utterly distinct, and do utterly different work. The two or more concepts can be held to be independently robust, with no entailment between the two in either direction posited, and indeed, there are extant views of this sort. I mention this because I will later examine one such view, and to say that there is nothing logically or metaphysically impossible about it.

However, these kinds of views are vulnerable to all of the problems I earlier identified. The important point is that the notion of priority is not meant to parse the *only* possible views, but instead, as a sort of heuristic, will indirectly give such views

an initial attraction. I hope I have shown that a view that posited two or more independently robust normative 'things' would have to demonstrate superior or unique explanatory power over and above 'first' type views, in line with the metaphysical burden it takes on.

Based on the above, we should be looking for a theory that, at minimum, meets the following constraints:

1. It respects the necessary grounding relation between the normative and the natural.
2. It tells us how normative properties, which are putatively causally inert, can prescribe to us, or even entice us.
3. It postulates one fundamental normative concept, or if it posits more than one, it demonstrates that the addition provides the theory with superior explanatory power, unique explanatory power, or both.
4. It offers a sufficient explanation of the concept(s) at hand such as it renders it/them sufficiently unmysterious.

In simple terms it should be consistent, clear, and parsimonious.

The goal of this chapter was to move from the attractive but vague notion that there is something robustly real about goodness or badness to a set of motivated desiderata that a philosophical theory aiming to give a more precise description ought to fulfil. In the next chapter I will compare and contrast three realist views about value according to this framework. The three views I will examine are:

Moore (1903), who argues that 'good' is the fundamental normative notion, with other normative notions (e.g., reasons) being entailed by the robust quality of goodness.

Scanlon (1998), who argues that 'reason' is the fundamental normative notion, with goodness being analysable in terms of reasons.

Dancy (2000), who maintains a 'parity' view of the sort discussed earlier: while Dancy agrees with Scanlon in the epistemic sense of *evaluation*, he partially sides with Moore in his metaphysics in that he maintains a distinct, independent and equally fundamental notion of *value*, that he argues does unique explanatory work.

Chapter Two : Moore and Scanlon

I take Moore's view first (1903). Moore's view both originates and sets the stage for just the specific literature I am surveying here; it is the most typical of the views at hand in that it approaches the question outlined above head-on and in the most obvious way: by taking goodness to be the obviously key concept when discussing goodness. Furthermore, Moore's answers to the second and third points mentioned in the introduction (what sort of values there are and how they are contextualised) are predicated on what he says value is in the first place (1903).

Moore's famous open question argument (OQA) (1903, 21) both inspires his particular view of value as a robust, non-natural property with certain specific characteristics, and inspires the current literature on non-natural properties in general (Feldman, 2005). Discussions about OQA of the second kind usually focus on its

putative ability to refute naturalism (Feldman, 2005, Frankena, 1939) as the 'naturalistic fallacy'.

While I wish to focus on what further shape OQA gives to the non-natural property beyond its status as such, we must still briefly consider the latter point: it partially inspires that shape. So, the first horn of the OQA is the one that justifies the nature of value as robust. It does this by first arguing for the negative conclusion that available arguments for a naturalistic view of value are wrong, before moving to the positive claim that non-natural arguments are more capable of explaining realist commitments. This is why it is sometimes referred to as identifying a 'naturalistic fallacy'. Recall my explanation of the apparently necessary grounding relation (Berker, 2018) between the normative and the natural. Moore is aware of this necessity and asks us to consider various natural properties that we seem to think entail goodness: for the reductive naturalist, the goodness is *just* in the properties (Moore, 1903, 10-14).

Moore points to the plurality of good things, such as that the list of natural facts or properties that can ground goodness is necessarily open-ended and never completely specified (1903, 43-44). Although, presumably, all the goodness in the universe must be grounded somewhere in the universe: Moore doesn't seem to suggest that there is unmoored, free-floating goodness that doesn't inhere in some actual good thing.

A competent speaker, then, can always ask a question of the form 'is pleasure good?' or 'is equality good?' on the understanding that the goodness must be something further than the pleasure or the equality: there are two distinct referents

being compared and contrasted (Moore, 1903). If goodness is reducible to pleasure, or any other natural kind or collection of natural kinds, then a question of the form 'is x good?', where x indicates the total of natural kinds that goodness reduces to, would indicate that the speaker is conceptually confused, or incompetent, in that the question would be of the form 'is x x?' (Moore, 1903).

But for Moore, any such question would always be justified until presumably x expanded to refer to the totality of natural kinds: at which point it would be both uninformative and a tautology: surely all the goodness is grounded somewhere. The putative fact that a competent and rational enquirer can ask a question like 'is pleasure good?' is meant to safeguard the notion that there is a distinct, independent and real 'good' that pleasure can be reviewed against: the question must be understood as being of the form 'is x y?'. The obvious plurality of good things that have only that goodness in common is precisely what indicates the type of goodness, of which they are tokens.

So while the open question inspires Moore's distinct view on not just the existence, but shape, of robust value, it is worth pointing out that it doesn't privilege Moore's view on the first point. This is because the same line of argument can be used to motivate other non-natural properties as such properties, such as rightness: we can analyse the question 'is x right?' in the same way.

Moore explicitly notes this and points out that his view isn't meant to motivate rightness (1903, 146-148). For Moore, rightness is not a robust normative kind but is instead independently extensionally empty in that it is completely specified by, and derived from, goodness. However, this view rests on a body of argument relating to

non-overlapping views of Moore's concerning moral duties and consequentialism (1903) which have been independently criticised (e.g., Ross, 1930).

I point this out because I will compare Moore's view with other robust ones that, when considered only in metaethical terms and leaving aside those ultimately independent points, are equally supported by OQA as far as robust properties go, and his view doesn't win by default. We will have to consider the rest of what Moore says about goodness.

The second horn of OQA tells us more specifically what goodness is beyond an abstract, robust kind. Apparently extremely unhelpfully, it tells us that goodness is indefinable (Moore, 1903, 8-10). This claim is meant to be a consequence of the above. Goodness cannot be defined in natural terms because of OQA; and cannot be defined as a derivation of other normative qualities like rightness. Since these are the only two options on the table, there is meant to be nowhere left to turn apart from towards goodness itself (Moore, 1903).

However, this point again seems to rely on first-order claims about consequentialism (Moore, 1903) and I don't assume that therefore no other second-order 'x-first' view can be true: this point is precisely what I am still investigating here. So instead, Moore holds that goodness can, at least directly, only be defined in its own terms, and this has the potential to be an unfortunate dead end. But instead, it marks the beginning of what could be called an 'epistemic turn' in Moore; one that inspires, and is echoed by, a similar turn in Scanlon and Dancy. So far Moore's definition tells us pretty narrowly what goodness 'is'. The key to further understanding this 'is', is to turn towards the obvious plurality of values, and evaluations, in the world.

For Moore, the fact that we have various distinct and apparently serious evaluative attitudes (e.g. aesthetic appreciation, love) towards apparently valuable things (e.g. paintings, people) is where we should really look. Because these things all seem different, yet the same, the key to understanding that one thing is to understand the commonality of these serious experiences. This is essentially a turn from an abstract to a practical definition of good. Never mind that goodness is 'indefinable' in the formal, philosophical sense: we can define good in practice as something that really seems to have that elusive quality and still continue to develop, through perception, our conception of the same (1903, 184-195).

To understand this, we have to briefly dip into intuitionist epistemology. There are at least three relevant ways we can use the term intuition when considering metaethics, and the Moorean intuitionist epistemology I am about to discuss seems to, in different places, incorporate all three. The first way this term might be applied is to a set of putatively common assumptions that contemporary philosophers have called 'folk metaethics' (Beebe, 2015): the philosophically undeveloped but otherwise reflective opinions of people about what ethical or evaluative views might be plausible. This is the sort of pre-theoretical assumption I referred to in the introduction, and an example I gave there was that a form of evaluative realism might be true (Darwall, 1998, Jackson, 2000).

The second way this term can be applied is to a specific view about theoretical, *a priori* certainty, in the sense that some truths are apparently 'self-evident' when apprehended. An obvious example of this is the grounding view: Hare's initial formulation (1964, 80-81) holds that considering the ramifications of grounding makes

the contradiction and confusion fundamentally and *intuitively* obvious. This is the opposite of the first sort of intuition: that one deals in common-sense acknowledgement of possibility, the second on reflective, rational acknowledgement of apparently obvious certainty.

The third way this term can be applied is in a practical, rather than theoretical, sense. Here, when we are actually in ethical situations, we can have direct apprehensions of goodness, badness, rightness, or wrongness. This view has recently been referred to as moral perception (Audi, 2010, 2013, 2015, Dancy, 2010) in the context of perceiving right action, and evaluative perception (Bergqvist and Cowan, 2018) when intuiting the presence of value specifically. When we are actually in the ethical situation, or perceive the good or bad thing as such, we can be said to have *direct* ethical or evaluative perceptions.

Indirect perceptions, on the other hand, are those where we aren't present in the situation, or directly perceive the object. An example that covers both sorts is when we read in a history book that the dodo became extinct in 1700 due to direct human action: a series of wrong acts leads to the loss of something valuable, but we have no direct access of any kind to either.

Moore's epistemology (whether theoretical or practical) takes one consistent approach to these considerations. He posits a 'moral faculty' that matches the non-natural realm: again because of OQA, but also because of our everyday experience, we can only perceive goodness directly or indirectly according to the perceptual view above: and we often can. On this view, goodness being indefinable is no impediment to finding out all about it: one has simply to be sensitive and to trust in the moral

faculty, whether in direct practical questions of the form 'is x good?' or in more abstract versions of the same.

Of course, Mackie is responding to precisely this literature and has something to say about all of this:

...If there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. ...if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else. These points were recognized by Moore when he spoke of non-natural qualities, and by the intuitionists in their talk about a 'faculty of moral intuition'... intuitionism merely makes unpalatably plain what other forms of objectivism wrap up... the suggestion that moral judgements are made or moral problems solved by just sitting down and having an ethical intuition is a travesty of actual moral thinking... however complex the real process, it will require (if it is to yield authoritatively prescriptive conclusions) some input of this distinctive sort, either premisses or forms of argument or both. When we ask the awkward question, how we can be aware of this authoritative prescriptivity, of the truth of these distinctively ethical premisses or of the cogency of this distinctively ethical pattern of reasoning, none of our ordinary accounts of sensory perception or introspection or the framing and confirming of explanatory hypotheses or inference or logical construction or conceptual analysis, or any combination of these, will provide a satisfactory answer; 'a special sort of intuition' is a lame answer.... (Mackie, 1977, 38-39)

However, Moore *does* argue for his view via OQA: the truth of OQA isn't simply assumed. What Moore actually does is use OQA to argue for his second-order view, where he incorporates claims that allow him to defend intuition on ultimately first-order concerns: whether particular x's are good, and how we might actually appreciate this. Moore's view is clearly intended to reach beyond 'sitting down' and philosophising into real aesthetic and ethical practice. For this reason, I think the problem is the plausibility of that faculty at that first-order level. Simply put, Mackie accuses Moore of defending the second kind of 'intuition' I outlined, while I read him as tending towards the third: at least when it comes to the actual practice of appreciating goodness. Moore's way of developing his definition of real goodness in terms of real love, real beauty, and apprehension of the same, is not *obviously* legislating from the armchair.

I point this out because it will be relevant later, but otherwise it seems the broader thrust of Mackie's objection succeeds (at least at highlighting the problem), when taken together with his criticism of robust properties themselves. A unique, metaphysical faculty that can sense unique abstracta is simply unmoored. The real problem here is that Moore tries to get out of one dead end with his epistemic turn, and ends up circling back to another.

To make this clear, let's take a (hopefully) obvious example of something good, and see whether it could have the causal powers Moore attributes to it. Since by ruling on a specific example we have to allow some ethical theory in, let us assume that Moore is actually also right about consequentialism: that goodness is what grounds

facts about correct ethical *actions*. Imagine that I have a dog, and that water is good for dogs; it is good for a dog to be hydrated.

This means that if, when my dog is thirsty, and it is the *right* thing for me to give him water, this must be because of the goodness of water for dogs. The abstract, Moorean property of goodness must not only be present (according to his consequentialism) but must be what inspires me to act. Now while my focus here is on evaluation, rather than action, that is precisely the point: consequentialism requires us to correctly apprehend and understand goodness in the first to maintain or maximise it.

I now turn to a further example. Imagine that (for whatever reason) we have a person who has had their memory interfered with in such a fundamental way that they do not know what mammals, let alone dogs, are, and they do not know what liquids, let alone water specifically, are. However, the Moorean moral faculty is completely intact, and the person is a perfect example of what Moore would call the morally sensitive person: before the procedure, they were perceived to be exceptionally attuned to picking up that goodness or badness.

This person is placed into a room with a thirsty dog and a bowl of water. Would they immediately sense the badness of an imminently dehydrated dog, the goodness of the bowl of water to the dog, and derive the right action of matching the two up? I think this seems unconvincing, and this is because it seems essential to grasp a range of ordinary facts to be able to make this judgement, including facts about mammals and facts about water.

The point here is that Moore's view does not correctly respect the ramifications of the grounding claim. Moore aims to explain the apparent plurality of values by pointing out what all values and evaluations are meant to have in common: that they involve goodness and responding to that goodness. However, what the plurality of values really have in common is that each one has some unique instantiation, and that the only way we can see to differentiate them is by their facts or properties.

Moore is correct that the broad variety of attitudes we have is relevant to defining value: but his epistemic turn ends up being a wrong turn. His solution to our apprehension of the plurality of value is an evaluative faculty, but when compared back to the simple, ultimately monist Moorean *property* of goodness, even accepting this faculty leaves us wanting.

To sum up: Moore defends his metaphysics from the thrust of Mackie's criticism via an epistemic turn, but this ultimately affirms that criticism. This isn't necessarily the fault of that epistemic turn, but of the interrelation of the two views. So, Moore needs to improve both sides of his view: he needs to tell us more about how the simple and indefinable property of goodness manifests in its plurality, and how we apprehend the range of that plurality.

And this is exactly what he did. The later Moore writes in *Ethics* (1912) that we do not in fact apprehend the simple property of value *directly* via a moral faculty: but instead that valuable things *give us reason* to value them. This is a new sort of claim for our discussion so far: while I have mentioned reasons in passing, I have waited to introduce them in context with Moore. So, what is a reason?

Just like realism, the notion of a reason has substantial pre-theoretical appeal as something that can explain normative notions. In day-to-day life, when discussing whether something is good or bad, we are used to asking others to tell us the *reason* why what they think is true, and offering our own reasons for what we think is true. When talking about the rightness or wrongness of actions, we usually prompt someone to explain their apparently troubling behaviour by asking what their reasons were for doing it, what made it actually acceptable to do; and when defending ourselves, we eagerly offer our own reasons. This ordinary way of doing things speaks to two theoretical notions discussed previously: that normativity is rational and logical, and that there is an intelligible relation between the natural and the normative. To elucidate this point with one simple example, consider the claim:

'I hit him because he hit me first'.

This statement argues from a cause to an effect: it seems to appeal to a rational view of normativity. Similarly, 'that he hit me' or 'that I hit him' are natural or base facts: they are meant to ground a further normative thing: that him hitting me was *wrong*, or perhaps that me retaliating was *right*.

Earlier I briefly discussed the idea of a normative category e.g. a distinct area of thought or practice that has in common with the others that it *is normative*. The two least controversial and widely discussed of these are goodness and rightness e.g. the 'shoulds' about what we should value or how we should act. More broadly, and more speculatively, we might hold that epistemology is normative in that it surveys what we *should* believe or reject, and how we should get there. While reasons are a paradigm example of a normative concept in that they point towards conclusions,

beliefs or attitudes that can be evaluative or ethical, does it follow that they are *distinct*? Perhaps in some way the term reason just refers to the constituents of already mooted normative categories.

The reason I point this out is that in all these areas it is natural to talk of reasons in a perhaps unexamined way. We can use the language of reasons to discuss the reasons we ought to value animal life, the reasons we ought to tell the truth, or the reasons we ought to accept mathematical or logical truths.

On this view, to say we *have reason* to think triangles have 3 sides is really to echo the observation of Moore's I quoted in chapter 1⁴: that the abstract notion of a triangle, when considered, is found to be rationally compelling. The point is that the facts just are the reasons: there is no separate thing called a 'reason'. Similarly, whatever way one defends the moral value of animals, the facts or concepts referred to will be the reasons to accept the conclusion. In this sense reason-talk really just refers back to normative talk about compellingness: that various kinds of facts impress themselves on us.

John Broome (2018) raises a related and reasonable issue about reason-talk: he agrees that we use the word reason broadly and ubiquitously, and that in *language* it serves an important, facilitatory normative role, but it doesn't follow that the thing referred to exists independently in the way discussed in chapter 1: e.g. that there is actually a particular robust property called 'reason'. For Broome, it's perfectly possible that the word 'reason' simply refers to the facts in virtue of which something is X, not the facts that ground something called a 'reason'. In this sense, reason-talk actually

⁴ Pages 18-19.

refers to apprehension of the grounding relation, of which facts are related to which others: to claim that natural fact X gives us reason to make evaluative claim Y is simply to propose that X grounds, or entails, Y.

This is directly problematic for what I am about to discuss. I am considering non-natural, robust views about value, and considering whether an 'x-first' type of view will be the most parsimonious while otherwise attractive. If the non-naturalist wants to say reasons are a distinct, unique, robust, normative property that can be used to analyse goodness or rightness, they first have to justify why they exist in a certain way.

The core consideration on this point in the literature is that reasons are a 'relational' property with certain characteristics. There are two different ways reasons are described as relational (Olson, 2018). The first way is that reasons are a third thing that relate on the one hand to some natural facts, and on the other hand to some 'thin' normative conclusion. For example, to say that any particular property *gives us reason* to think that the thing is bad; it recommends the conclusion. The relational view of reasons tells us both where they come from, what they go on to do, and why they are distinct from the other relata. If the orthodox relational view is correct, reasons are a distinct normative kind over and above any particular normative place they might weigh.

Olson (Olson, 2018) draws out a distinction between two ways we use the term reason that helps explain why they are a normative kind: a distinction between explanatory reasons and normative reasons proper. So far, I have referred to both sorts without clearly drawing this distinction. Explanatory reasons are those that

explain some phenomena, whether moral, evaluative or not, without necessarily *justifying* the thing. Olson uses examples such as facts about planets, or about dinosaurs, as explaining those things to us without having a normative, justificatory or binding sense.

When we discuss a normative or binding reason, however, we absolutely expect compellingness to feature in our definition. These reasons don't have to be ethical: again, they can recommend different sorts of compelling conclusions. Alternatively, explanatory reasons certainly can be evaluative, epistemic etc. insofar as they *explain* why we hold certain beliefs or attitudes. The point here is that morality doesn't have a monopoly on justification, and other areas of discussion do not have a monopoly on explanation: defending ethical reasons does not mean defending anything completely special or unique.

But focusing on the metaphysics of reasons doesn't allow us to address Mackie's complaints above directly. To do that, we have to combine the relational, grounded view of reasons with the epistemic potential of reasons. It is because reasons are relational that they seem to be able to explain the plurality of values and evaluations: but why should we hold that moral or evaluative reasons are any less mysterious than the Moorean property?

The epistemic potential of reasons is why Moore allows them into his ontology. The Moorean property of goodness is apprehended by a mysterious moral faculty; reasons are apprehended by a familiar rational faculty. We can take a companions-in-guilt type approach (Lillehammer, 2019) that Mackie's criticism weighs much less heavily on. This is because we are already happy to accept the compellingness of

reasons in other familiar areas: that there is a particular phenomenological ‘feel’ to apprehending that x is a reason to accept y. These could be mathematical or epistemic statements.

Moreover, reasons are *specific* and *local* in a way that the thin, universal property of goodness simply isn’t. Since reasons are relational, always a *reason for*, they have no trouble explaining the open-ended plurality of good things and aesthetic, evaluative or ethical attitudes. In my example above, that facts about dogs or water give us *reason* to draw certain conclusions, this is a completely familiar and uncontroversial epistemic claim that seems to be as acceptable as saying that facts give us reasons to draw certain conclusions outside of the ethical or axiological sphere. There is no need for a special, unique value quality or a correspondingly special faculty: both of these only serve to explain limited phenomena at high profligacy costs, whereas reasons are both easier to explain and use to explain. Again, Mackie criticises robust realism, or indeed any realism, on the following practical epistemic grounds:

“...however complex the real process, it will require (if it is to yield authoritatively prescriptive conclusions) some input of this distinctive sort, either premisses or forms of argument or both. When we ask the awkward question, how we can be aware of this authoritative prescriptivity, of the truth of these distinctively ethical premisses or of the cogency of this distinctively ethical pattern of reasoning, none of our ordinary accounts of sensory perception or introspection or the framing and confirming of explanatory hypotheses or inference or logical construction or conceptual analysis, or any combination of these, will provide a satisfactory answer; ‘a special sort of intuition’ is a lame answer...” (1977, 38-39).

The solution I am defending is that there is no distinct ethical sort of *reasoning*, and that defending Olson's approach above allows us to see the important similarity: what explanatory and justificatory reasons have in common is that they are reasons, and apprehended by *reason*. Again, there can be putatively explanatory reasons that are either ethical or not, and normative reasons that are ethical or not: reasons can be more broadly normative in that they are epistemic reasons, for example. However, this line of thinking simply does accord with our usual accounts of reasoning and argument, contrary to Mackie's criticism.

So reasons are a sensible and useful addition to Moore's ontology, and ought to protect him from Mackie's criticism, and indeed, they do: insofar as telling us what reasons go on to do. But as to where reasons come from, Moore maintains that reasons are grounded in goodness itself: the singular thin property of goodness is what inspires all those different reasons we have. Instead of immediately criticising this differently, I will first turn to some over views about reasons in a way that will highlight why this is ultimately a troubling 'improvement'.oqa

Thomas Scanlon (1998) has recently developed a model of value that relies on, and respects, several of the key notions we have discussed: specifically the predication of the normative and the natural, the attraction of parsimony, Moore's OQA about a robust notion of goodness, the normative primacy of reasons, and the popular relational conception of reasons.

Moore's open question led him to conclude that while the list of what kind of things can be good in what places is necessarily open ended, this plurality of good things has in common the possession of the master property of goodness. Mindful of

the problems Moore ran into with this approach, Scanlon agrees that the plurality of particular good things may be open ended, and that the open question is a valid one, but instead points out something else all good things have in common: that facts about, or properties of them, always give us reason to value them in the way they are valuable.

This approach purports to give a similarly complete description of what it means to be valuable, while avoiding the problems Moore had with causal impotence. This point is simply an observation of the outcome of respecting the grounding claim, that Moore also incorporates: There is nothing directly novel about Scanlon's point that the facts that ground the putative goodness are a good place to focus when trying to escape Moore's epistemic dead end.

Following from this, Scanlon concludes that the separate, distinct and robust property of goodness is unparsimonious, extraneous and superfluous, and this is what is novel about the way Scanlon's discussion progresses. If good things seem to be good in virtue of the actual properties they have, and these properties completely explain how we come to apprehend that goodness, there seems to be no obvious metaphysical reason to include a further property of goodness, as opposed to a *response* to those properties that the thing that has them should be viewed positively.

This view is supported by considering two problems with keeping that property in our ontology: an initial point about plurality, and a derivative about superfluousness. The problem of plurality is that there seem to be a necessarily unlimited amount of reasons. A complete reason to X necessarily refers to all sorts of other contingent features, including the features of the object at hand. Since for Moore there is just

one simple, 'thin', property, there is no room for all the reasons there will ever be to value something to be inspired by that property alone, with no reference at all to any ordinary features of objects or situations. Consider my dog/water example above: it seems strange that the same single property could tell us precisely what is good about each thing without referring to the facts that give them their apparently unique shape.

The problem of superfluosity is outlined in the literature by Jonathan Dancy (2000, 162-165), where he makes a point that builds upon and strengthens the one above. For Dancy, not only does goodness at minimum have to share with all kinds of ordinary features in terms of filling up the specification of a reason, but there is just no space for goodness at all when it comes to reasons: reasons are completely explained in relation to natural features.

Dancy uses the example of a toothache (2000, 164). All the natural facts about the toothache: that it is painful, potentially harmful to health, will interfere with other pursuits, etc. together give us reasons to understand that it is bad. Having done so, it simply makes no sense to say 'this realisation that it is bad gives me a further, extra reason to think it is bad'. The badness of the thing cannot itself give us reason to think it so, as Moore's second (1912) view would have it. What this point really serves to do is to show that Moore's second attempt simply passes the problem back a step: the reasons Moore's position struggles to explain how goodness causes our attitudes are similar to those his second approach struggles to explain how it gives us reason to adopt those attitudes: his view that we have reason to appreciate good things is true precisely because it is tautologous.

Note that Scanlon and Dancy unite in holding that reasons are indefinable on similar grounds to Moore with goodness: as robust relational properties they are irreducible to their relata and therefore cannot be defined *directly* in any other terms apart from reason terms. Just as 'good things are good' is true but empty, so is 'reasons are reasons'.

However, reasons have a clear advantage in that they have resources to respond that are not available to goodness. Just as Moore attempted to give a practical definition of goodness in terms of our apprehending the presence of it, Dancy gives a practical definition of reasons on much the same lines (2004b, 142). This move works where Moore's doesn't because reasons have none of the causal impotence of Moorean goodness; it isn't reasons that give us reasons, but ordinary, accessible facts and properties that do.

While we hold that reasons impress a particular conclusion on us, this is a different thing to reasons telling us that they are reasons at all. While Moore refuses to allow anything apart from goodness to either compel us or give us reasons that compel us, Scanlon and Dancy are happy to allow that ordinary facts and properties do this, and since this claim is easier to accept, its use in supporting a working definition of reasons is easier to accept.

So, Scanlon's view holds that for an object to be valuable is for it to possess facts or properties that give us reasons to conclude that it is such. It 'passes the explanatory buck' from the Moorean property of goodness to the similarly robust, but more accessible property of being a reason: this is why he calls it a buck-passing

account, and in the context of evaluation, it can be referred to more specifically as evaluative buck passing (EBP).

The first point to make about the buck-passing view for our purposes is that a close readings shows that it refers to two different sorts of reasons, reasons to have certain attitudes, and reasons to carry out certain actions. Scanlon refers to reasons to defend, preserve or promote; and reasons to admire. The first sort clearly refer to taking some further action, while the second simply refer to apprehending the value of something. I concur with Phillip Stratton-Lake (2013, 79) that the view is (at least for our purposes) weakened by including both sorts of reasons, and that a clearer and stronger version only weighs on attitudes. This is for two reasons: because it is irrelevant, and because it incurs extra explanatory debt that (because of the irrelevance) is simply unnecessary.

It is irrelevant because we are asking specifically is what it means to be good, or to know that something is good: not how we should act in response to it. We would have to be confident about our answers to the first sort of questions before we could confidently say how we should go on to act. It incurs extra explanatory debt because discussions of right action involve consideration of all sorts of other factors: facts about the agent say, or about the situation, outside of any good objects. But this point depends on arguing for an answer to a certain question: whether there can be purely 'attitudinal' reasons that have nothing to do with any possible action.

It seems that to prove that evaluative reasons are distinct, we would simply need persuasive examples of such reasons, and these are easy to find. Consider any situation where there is clearly something valuable at stake, but nothing one can do.

A popular example (Singer, 1972, 231) of a clearly undemanding moral duty is to imagine an adult who finds a small child struggling in water that is shallow for an adult: clearly the child being in trouble is bad and the cost of helping (getting one's shoes wet) is held to be trivial. However, what if the adult is paralysed and bound to a wheelchair? It seems they still have reason to *evaluate* the child's trouble exactly the same, to obviously recognise that it is bad, but no reason at all to wade into the water. This last point is crucial: the relational view of reasons implies that for there to be reasons *at all*, there have to be relevant facts.

Discussion of reasons to *act* requires consideration of facts about agents, for example, outside the thing we are saying is good or bad. But distinguishing purely *attitudinal* reasons allows us to avoid being drawn off on a tangent: we want to know whether the thing is good or bad, and therefore that the agent should think it so: not whether they should go on to do some extra, distinct thing inspired by the acquisition that attitude.

A similar example is to consider places where there is again *no reason to act at all*, but because of facts about the situation, rather than the person, as in the above focus. Again, Scanlon mentions reasons to 'preserve' as well as reasons to 'admire'. Preserving is certainly a form of action: preserving a beautiful painting certainly involves acting, and acting in light of some good.

Consider a tour party visiting a museum. In the museum, all of the objects are carefully tended by experts, kept in temperature, humidity and light controlled environments, and heavily guarded. The tour party will be given reasons to admire or appreciate the relics, but arguably reasons to preserve or defend will simply not exist

for them in the first place: these things have already been completely and competently done. Again, we can distinguish a narrower conversation about attitudes, and a broader one that involves attitudes and actions. The conclusion to be driven home here is that when analysing value and evaluation, it is an optional extra to go on to discuss related action, and will require intervening in many other debates that do not touch directly or even indirectly on the meaning of 'valuable'.

The second point is that the buck passing view can be divided into positive and negative theses. The negative claim is that for an object to be of value is *just* for it to possess certain properties: it is impossible for value itself to be reason providing. The positive claim is that for an object to be valuable is for it to have properties that inspire a certain positive attitude towards it. The negative claim reiterates the problem with Moore's second view, which we already examined, while the first is meant to present an attractive and parsimonious alternative according to the responses to that view. In the next section I will separately discuss how critics have, or have not, problematised the positive and negative views respectively.

Chapter Three : Dancy and Scanlon

So far, I have briefly discussed the initial attraction of the buck-passing view and what makes it superior to Moore's view: the value first view. But I have also mentioned the context I make this comparison within: Moore and Dancy predicate their views about contextuality on their prior views about value, and accepting the buck-passing view would have ramifications for the form of these other views. So,

before I turn to addressing Dancy's criticism, I should also want to be able to say why the view ought to be attractive specifically to Dancy.

And the reason for this is that Dancy endorses core parts of Scanlon's programme (the reasons first view) about the rightness and wrongness of actions and derives a structurally identical model of those concepts: the deontic buck-passing view (Dancy, 2000, 2004b, Metz, 2020). Dancy thinks both that facts about actions, situations and agents give us reasons to find those actions right or wrong, and that this is all there is to it: rightness or wrongness can never provide reasons.

Accepting EBP would allow Dancy to take a consistent 'reasons first' approach to all kinds of normative talk, as we have previously mooted. The only reason to prefer a different, less parsimonious approach would be if that approach could add something unique in terms of explanatory power, or alternatively, point to a deficiency in the explanatory power in EBP. And this is exactly what Dancy does: he claims that EBP is overly parsimonious to the point that gaps appear in its explanatory ambitions: there is still specific, unique work for the separate property of value to do. The point that Dancy endorses deontic buck-passing only has any relevance if the thin properties of goodness and rightness are alike on key points.

For this reason, Dancy walks a middle path. While agreeing with Scanlon about *evaluation*, he stays closer to Moore on *being of value*. Dancy claims that values and reasons share the same grounds (the features of the valuable object). So what does Dancy say we still need a specific, distinct value property to explain? Earlier I pointed out that EBP has at least four different dimensions: the positive and negative claims about actions and attitudes, respectively. Dancy is happy to accept all of the negative

claims made: that goodness itself, as an independently robust property, cannot ground reasons to have an attitude, or to act, and this is congruent with his objections to Moore's views on the same.

Instead, he focuses his objections on the positive claims: that there is no such independently robust property, and that to be valuable is instead to *just* possess certain reason-giving facts or properties. Since I agreed with Dancy, Stratton-Lake and others that buck-passing about *actions* is indeed a problematic position, and that I instead intended to promote the view only about attitudes, I will only consider Dancy's criticism of the same.

While Dancy's text contains four complaints, only two of them apply to the positive claim about attitudes: the other two problematise EBP's ability to completely explain right action. Both objections have been influential and spawned several further versions in the literature, so those are the ones I will discuss. The first of these is the observer objection. The objection, simply put, is the view that there is a subtle but important difference between the concepts of value and of evaluation: they are non-overlapping: specifically, something can be *of value* without being *valued*.

Andrew Reisner formalises this as values and evaluative reasons having 'unlike variance conditions' (2021, 3) and identifies several later versions of the point in the literature that he corrals under the heading of 'under-generation' arguments (Bykvist, 2009, Dancy, 2000, Heathwood, 2008) in that Scanlon's model 'under-generates' value where we otherwise expect to find it. This approach is also often referred to as the wrong kind of values problem (WKV).

Dancy's original version of the point (2000a, 171) asks us to consider a person experiencing great hardship, far away from any other people, and to decide whether this would still be bad according to EBP. According to Dancy, it would not: EBP tells us that value is a result of shared or objective reasons to value the thing, and in this case there are no shared reasons, and therefore, no value. This point depends directly on the view that reasons are a relational property: they require the existence of both relata (in this case, the thing that is good or bad, and the rational and appropriately sensitive observer) to be instantiated. Dancy expresses the point in simple terms:

...reasons belong to, are for, individuals. There are no reasons hanging around waiting for someone to have them...But goodness is not like this. Something can be good or bad without specification of an agent... (2000, 170).

This sort of objection, then, aims to demonstrate that EBP violates strong intuitions we have about a realist picture of value: that goodness or badness are enduring, and that they are carried about by the things they belong to. The theory is held to be defective in that it 'under-generates' goodness according to what we seem to want it to do. If successful, this point would demonstrate that EBP fails one of the tests I claimed plausible realist views must meet: it fails to analyse some part of normativity and therefore has inferior explanatory power to competing theories.

The first thing to note about Dancy's version of the point is that it requires some modification, in that the person suffering great hardship would be the required observer. Philip Stratton-Lake (2013, 93), working with the kind of rational, reason-based approach to evaluation we have discussed, asks us to instead consider a non-rational animal suffering alone, one incapable of apprehending reasons, and this

seems to preserve and clarify the point perfectly well whilst avoiding the issue. Krister Bykvist (2009, 15) offers a positive version involving a possible world containing only *happy* egrets, and asks us to consider if there can be any reasons or attitudes in this world. According to Bykvist, the relational predication of EBP makes it questionable whether this completely happy world is good, which seems to be an otherwise obvious conclusion.

Naturally, responses to the objection have also appeared in the subsequent literature. Stratton-Lake (2013, 94) and Jussi Suikkanen (2005, 2009) offer distinct defences of EBP regarding WKV. Stratton-Lake offers a 'companions in guilt' type strategy, where he points out that the objection problematises not just Scanlon's view about goodness, but the views of Dancy and others about rightness: notably the deontic buck-passing view (Dancy, 2000, 2004, Metz, 2018) described above. Unfortunately, I believe this strategy is incomplete insofar as Dancy holds a purely verdictive conception of rightness and wrongness (2000, 166) in that they are instantiated when someone has reasons pertaining to them, and exist contingently in relation to actual situations and people.

Dancy points out correctly that we sometimes talk about goodness as if it is a purely verdictive concept: that talk about goodness is talk about us *passing verdict* on whether things are good and bad. But for Dancy, this inappropriately reduces *being of value* to *being evaluated*. Really there are two ways the term goodness applies: a verdictive sense and an enduring one, where goodness inheres something, is 'carried about' by that thing, as in Dancy's example, specifically without the specification of an agent. Again, Dancy is happy to agree with Scanlon on how we actually apprehend

value: just not on what its existential status is. So Stratton-Lake's position only survives if goodness is not importantly different on this point about endurance: this is what would make EBP and its deontic cousin 'companions'.

A second, and more promising, way of defending the view then is to point out directly that Dancy's criticism seems to rely on a misinterpretation of Scanlon on whether his model is purely verdictive. Dancy and Bykvist correctly point out the relational nature of EBP, but predicate the details of their objections on the existence of reasons or attitudes respectively, whereas Scanlon's version is predicated on *facts*. The criticism, then, would weigh more relevantly on a purely attitudinally predicated approach rather than Scanlon's. While reasons or attitudes can be argued to be contingent upon particular situations with particular observers, and hence, completely explained by a verdictive goodness of model, facts about things, and properties of those things, presumably *are* 'carried' about by them.

Jussi Suikkanen has made an argument on these lines (2005, 2009). He points out that Scanlon indeed predicates value on facts or features, not on attitudes or reasons. He writes:

...Dancy appears to have thought that in order for something to have reason providing properties the necessary condition is that it gives actual reasons... There is a clear fault in Dancy's argument. It seems odd to think that in order for something to have reason-providing features, it must actually give someone real reasons. Consider an analogy; a fully functioning loudspeaker in outer space has sound-providing features, even though it cannot provide sound in that situation... (2005, 25).

This analogy is a useful one in that it highlights that crucial distinction we keep coming back to: the difference between the phenomenal experience of reasons (or sound) and the real facts that ground that experience. As Scanlon writes:

...to call something valuable is to say that it has other properties that provide reasons for behaving in certain ways with regard to it... (1998, 96).

The absolutely crucial phrase here is 'it has other properties'. Scanlon indubitably predicates his argument on the existence of facts or properties, and makes reasons and attitudes contingent on them, in that order. And since those facts or properties exist just as well in Dancy's example as any other, I think we have to rule in Scanlon's favour on this point.

Accepting this plausible resolution allows us to return to the first, indirect strategy on slightly different lines. Having established the safeguarding role the focus on facts plays, and that Scanlon's view isn't a purely verdictive one, we can apply a companions type strategy (Lillehammer, 2019) but with a focus on existential status as well as epistemic considerations.

Consider again the putative similarity between mathematical objects and judgements, or epistemic ones, with evaluative ones. Those areas similarly involve a reason-relation of the kind Scanlon defends: a grounds-reasons-agents relation. Presumably the notion of a triangle exists whether anyone has apprehended it or not, just as truths are discovered rather than invented. Accepting Suikkanen's approach means accepting there is no special metaphysical problem here: we can confidently accept a companions type approach (Lillehammer, 2019) not just on attitudinal or epistemic, but existential grounds.

So, Stratton-Lake defends EBP indirectly while Suikkanen offers a direct defence based on the view itself: I argue they should be accepted together, and that Suikkanen's is more directly relevant, both because it lends support to Stratton-Lake's and because it is also the one Dancy pre-empts as a plausible reply: it actually inspires his *second* criticism. This second objection instead criticises EBP for *over*-generating value, and in contrast to WKV above, is instead known as the wrong kind of reasons (WKR) problem in that the reasons present are *not* ones that have anything to do with the putative value of the object itself. Dancy writes:

...To have potentially reason-giving features is a less polyadic matter than to have actually reason-giving features, since something could be of the former sort without our needing to specify any particular individual for whom the reasons are reasons (since the reasons don't yet exist). This manoeuvre seems sound, so far as that goes, but it raises worries of another sort. For it seems far too easy to have features that are potentially reason-giving - that would, in certain circumstances, give us reasons. Something that has no value at all might well have features that would, in certain circumstances, ground reasons... (Dancy, 2000, 170).

This passage again invokes the relational status of reasons and accepts that EBP focuses on the first relatum (natural facts). The problem, then, is that the same feature of EBP that potentially protects it from the argument from under-generation, is precisely the one that makes it vulnerable to overgeneration. The problem is now precisely the reverse: something can be *valued* without being *of value*.

EBP can over-generate value by recommending we value something that has no value at all, which violates another obvious intuition about value: that if we are

realists about value, and we think binary claims of the sort “x is/is not valuable” have truth values, we should only value things that are actually valuable, according to the *way* they are valuable. This is part of a broader worry about ‘fittingness’, which is the notion that our reasons to value something, together with the attitude they recommend, should be fitting to the thing.

To make this point clearer by example, consider a truly beautiful painting, and two people who claim to value it immensely. However, when questioned, the first person claims to love the painting only for its aesthetic qualities, while the second confesses that they appreciate the painting only for its enormous monetary value. We would normally want to say that the first attitude is simply more ‘fitting’ to the object at hand, and more than this, that the sort of independent, objective reasons that Scanlon talks about will be fitting to the object: they will recommend the first attitude and not the second. Certain sets of facts seem to not just call for a positive response, but a specified positive response.

While we might otherwise want to claim that the painting has the specifically evaluative property of *beauty*, for example, and that it is such a specific property that matches up with the correct response, accepting EBP means rejecting talk of such specific non-natural properties, and instead claiming only that it must be specific natural facts that give us the robust reasons for such specifically correct conclusions.

Cases like this might not be especially problematic: we could hold, for example, that the attitude of the second person *is* grounded somewhere: just in economic facts, not aesthetic ones. If we take this kind of view, it doesn’t seem that EBP validates the relevant reasons and responses at hand: it tells us that the painting is valuable in

relation to attitudes like that of person one, and that is what is relevant to EBP as a definition.

The first person to combine general worries about 'fit' with specific criticism of EBP in the literature in a concrete and genuinely problematic way was Roger Crisp (2000, 459). Crisp asks us to imagine an evil demon who commands us to worship a saucer of mud for its own sake. If disobeyed, he will severely punish us. The specific point here is that the demon can tell if we actually admire the object or not, and whether we do so *for its own sake*. For brevity, this evil demon example will from now onwards be referred to as ED.

This is why my suggestion about how we might approach the painting case above doesn't obviously hold. However tempting it is to simply claim that we are really talking about the badness of severe punishment, or rather the goodness of not experiencing such, and our correct appreciation of the same, ED demands that the facts must be about the mud alone, and the attitude towards the mud alone. This is meant to be the problem with predicating EBP on facts about *specific* things rather than other places in the reason relation: when we find a scenario where there are indubitably reasons, but no relevant and *suitable* facts about the object to ground those reasons, EBP again seems deficient in explanatory power.

While there are several ways we want and need to reply to this objection, the first point to make is that Scanlon clearly considered these issues, and at least *attempted* to build resistance to them into his original formulation in two key ways. The first of these is that Scanlon is a pluralist about the kinds of evaluative attitudes we can have, he holds that there will be as many different ways of valuing things as

there are distinct sorts of things, and that we will value them according to the way they are. Scanlon clearly points out that the attitude (as in my painting example above) will be one that suits the thing being admired: it must have the right sort of attributes to generate a certain attitude. The key passage is this:

...Exactly what these reasons are, and what actions and attitudes they support, will be different in different cases. They generally include, as a common core, reasons for admiring the thing and for respecting it, although 'respecting' can involve quite different things in different cases... (Scanlon, 1998, 95).

A second, and related point is that Scanlon tells us that something is truly valuable in the realist sense if others have reason to value it in the way we do: the reasons must be mind-independent and universal. This isn't meant to preclude the possibility of real reasons that don't meet these standards: it instead focuses on the point at hand, which is things we think are independently and robustly valuable. This point implies the first: if there is one set of shared, consistent reasons, there will be a particular attitude specified for everyone, *ceteris paribus*, to take. Again, there is a key illuminating passage:

To value something is to take oneself to have reasons for holding certain positive attitudes toward it ... To say that something is valuable is to say that others also have reason to value it, as you do (Scanlon, 1998, 95).

Stratton-Lake (2013, 82-83) points out that this may map onto another common and putative distinction: that between an agent-relative and agent-neutral reason, and that buck-passers may or may not endorse such relative reasons. This is a formalisation of a point I just mooted: the notion of real reasons that are

nevertheless not universal or necessarily robust. Examples include the reasons we have to specially value our own dog or our own wedding ring: these are not universal reasons for everyone, but it seems we do have reasons to value them and that, at least for us, they *have value*.

While I will return to this later, the point here is that I agree that the two distinctions may or may not be either non-overlapping or co-extensive: what is relevant is that agent-neutral reasons are the ones referred to by Scanlon and will be the important kind. Whether or not agent-relative reasons exist, there would be no problem for over-generation there: presumably how many agent-relative reasons there can be is completely contextual or contingent. There will only be a problem if universal, agent-neutral reasons are generated, yet according to EBP the other two places in the reason relation are respectively unfitting, or absent.

The first response to the problem in the literature is from Rabinowicz and Ronnow-Rasmussen (2004), which defends EBP on the basis of a distinction drawn between two places reasons can be grounded: object-given reasons and attitude-given reasons. A similar distinction is elucidated by Christian Piller (2006). Rabinowicz and Ronnow-Rasmussen point out that mental states may be natural kinds that are suitable grounds for reasons: and that the mental state of desiring the mud, rather than the mud itself, may be the *object* of evaluation. While they cast doubt on whether their own response really stands up, Philip Cook (2008) Gerald Lang (2008) and Stratton-Lake (2005) have all defended the spirit, if not the letter, of the approach, and offered more recent glosses.

I think defenders of EBP are slightly missing the point here insofar as they focus on *where the reasons are grounded*. It is natural to think in this way precisely because of the approach to solving WKV (Suikkanen, 2009) that I endorsed: EBP focuses specifically on those real properties which are meant to ground reasons, and the apparent problem that ED creates for EBP is that there are no real grounds, yet apparently real reasons.

However, this isn't a special problem for EBP at all, but instead for any view (including Dancy's) that holds that evaluative epistemology is about reasoning, and that there are distinct reasons to hold certain attitudes: the kind of views I have promoted in this essay. This is because ED would be a problem for these views whatever the grounds are held to be. Consider Moore's view again: that the robust property of goodness gives us reason to appreciate objects that possess it.

Presumably a Moorean would hold that the saucer of mud does not possess this special property: yet ED tells us there are reasons to value it. ED is thus equally a problem wherever the reasons are meant to come from: it just tells us that the necessary entailment isn't present. It seems that whether we say goodness, objects, or attitudes ground the reason, the ED defender can simply dig their heels in. This is the wrong kind of *reasons* problem precisely because it is the reasons specifically that are problematic. So, really we need an approach that tells us why there are no true evaluative reasons hanging around without suitable moorings.

Danielson and Olson (2007, 517) recognise the point just made: that ED isn't a special problem for EBP but instead for all kinds of moral and evaluative realists. Whatever definition of goodness and reasons one offers, ED severs the apparently

necessary relation between the two. They aim to improve on the dual reasons strategy, while endorsing the spirit of it, by moving from a discussion about dual grounding places to reasons alone. They write that EBP should be qualified to include the claim that there are two sorts of reasons, and that it is only when both sorts are present that the analysis is relevant.

In doing so, they attempt to incorporate the 'fittingness' clause that I claimed is an important positive feature of EBP, but in a way that more clearly avoids the ramifications of ED. It is tempting to directly claim that admiring a saucer of mud would be unfitting, and that EBP doesn't generate *either* real reasons or value here. But this is why we require the supernatural power of the demon: in telling us that we must admire the mud for its own sake, he seems to suggest that to survive, we must *fit* our attitude to the mud, or rather, somehow fit the mud to our attitude.

Their solution is to reason along the lines that fittingness is itself a normative category we can reason about and have distinctive reasons for. Just as I have claimed there are evaluative reasons and reasons for action, Olson thinks there are reasons to find our results in any of these areas fitting or unfitting: a sort of metanormative reasoning.

On this view, then, we have two sorts of reasons: content reasons and holding reasons. Holding reasons, as the name implies, are simply real reasons to *hold* a particular attitude, whether or not those reasons are inspired by the features of the object at hand or not. A relevant example is, in ED, the fact that the demon will punish us seems to give us a real reason to hold the relevant attitude of admiration, although this is *not* a fact about the saucer of mud, which is the relevant object at hand. Content

reasons, on the other hand, are meta-reasons to find the fit between attitude and object correct.

Olson concludes in a later paper that the benefit of having some particular attitude *cannot* contribute such a content reason. I understand this to be because of priority of logical explanation: EBP tells us that the attitude has to be fitting to the object, and this will be via the reasons. If one is considering whether some attitude is fitting or not, one cannot use that attitude as evidence of such. Otherwise, we could conclude that 'a has a reason to hold attitude x towards object y because it is fitting to do so, and it is fitting to do so because x is a fitting attitude'. As Olson points out, EBP extends an object-focused ontology, not an attitudinal or reasoned one:

...But, one might reasonably ask, why are [attitude]-references within the [fittingness]-clause illegitimate? The answer seems to me fairly straightforward: what we are interested in is value analysis, that is, the formal question of what it is for an object to be valuable. We are not interested in what kinds of attitudes it would be useful, harmful or intrinsically worthwhile to take up towards various objects. So the restriction imposed on the content of the [fittingness]-clause is justified by the very purpose and object of our endeavour: it is not just an ad hoc manoeuvre introduced in order to escape the WKR objection... (2009, 300).

However, it was necessary to focus on reasons and attitudes precisely because this is what ED does. While Olson's solution, like those of the first group above, and those we surveyed when responding to WKV, ultimately refers back to the specific formulation of EBP, it does so by accepting that ED doesn't allow us to defend EBP so directly: however we claim our reasons are grounded,

This second type of approach is more promising in that it focuses directly on the real problem. Moorean and Scanlonian approaches alike tell us that wherever there is real value, there are real, substantial reasons to appreciate it. ED tells us that such reasons can exist without that qualifier: this is the precise issue. Olson problematises whether they are such substantial reasons: and I think his approach is ultimately the most attractive in that it draws on the *fittingly* relational status of EBP as a positive, when it is apparently what instantiates the problem.

The goal of the previous two chapters has been to survey which approaches to explaining robust realism about normativity are the most promising, and to conclude that 'reasoned' approaches seem to most attractive according to certain desiderata. I considered the two main ways a separate conception of a value property may still be required, and concluded that neither succeeds in problematising EBP. This concludes my discussion of the first of three questions I mooted in the introduction: what it means to be valuable *at all*. In the next section I will move to consider the second of those questions: whether there are distinct *types* of valuable things inside this picture.

Chapter Four : Intrinsic and Default Values

Now we have considered what 'valuable' might mean in the first place, we can move onto our second and third questions: whether there are categories of good things, or whether they can change in their evaluative status. I will take the second question here, but will refer to the third so far as answers to the second weigh there.

This is because the answer to the second question may partially answer the third, in that one of the evaluative categories may be contextually contingent by default.

Recall the three ways I identified that values might vary: in status, in valence, and in amount. Something can become more or less valuable, turn from good to bad, or turn from good or bad to indifferent. However, there is a different sort of variation; one that doesn't speak to contextual or holistic views but is broadly accepted across the philosophical spectrum. This is the notion that there are two or more distinct types of value. For example, in one sense we think that differences in value are just differences in amounts of the same thing. For example, we would presumably hold that £10 is worth more than £5, or that two paperclips are worth more than one paperclip.

An apparently similar proposition would be 'human lives are worth more than paperclips'. It is easy to mistake this at least partially as just representing a very significant difference in *amount* of value. On this view, the sentence 'human lives are *more* valuable than paperclips' and the sentence 'two paperclips are *more* valuable than one paperclip' differ only in the size of the gap, and are otherwise alike.

But this isn't how we are used to using these concepts at all. Instead, we seem to think that human lives aren't just worth more; they have a special, very serious sort of value: such as they are always worth more than any amount of paperclips. So far, for simplicity, I have blurred over this distinction and used examples that are both trivial and non-trivial, putatively ethical and apparently nonethical, but reflection on ordinary expression and intuition shows a tacit understanding of what philosophers might call the conditioned or unconditioned nature of value.

Investigating this distinction is crucial because the Moorean and Dancyean views on our other questions completely rely on it. Again: it seems that value can vary in evaluative status, in valence, and in amount. And it seems possible that all of these points could be completely explained by appealing to the spirit of the consideration above: that there will be some few special and important things that have fixed, invariant value, but a myriad of contingently valuable things that, when we call 'good', are simply supporting some real value in some way. We think the first sort of thing can change in all those ways, while the second sort doesn't seem to at all.

But Dancy and Moore don't think this: they think both that there are things that have a special form of value, but that even those things, either alone or together, can change in contextual ways, and that is what makes these views interesting and somewhat controversial. So for Moore and Dancy to say something interesting and non-trivial, they need to tell us how something that *seems* stubbornly fixed, special or invariant can vary. In other words, they have to argue that there are at least two sorts of value. If there is only one sort of value, contingent value, there is nothing for these views to explain: their commitment to rebutting static approaches requires at least a partial endorsement of those approaches to be argued against. So how do philosophers conceptualise value categories?

Theoretically, Philosophers have coalesced around two main ways of parsing values: the final/instrumental distinction (Koorsgaard, 1983), and the intrinsic/extrinsic distinction (Moore, 1903). The first of these focuses on our relationship with and apprehension of value: an epistemic approach that begins with the rational, evaluating subject. The second focuses directly on putatively valuable

objects in isolation: a metaphysical, ontological approach. Due to my endorsement of a buck-passing definition of value *per se*, and specifically one that focuses on grounding facts, I will mainly consider the ontological distinction here.

So, the intrinsic/extrinsic distinction, instead of focusing on how we value things, asks us to consider the properties of the putatively valuable things directly: our question seems to be an existential one about two or more kinds of things in the world, and perhaps this is the most forensic way of doing things. Intrinsic values, as the name suggests, are those things whose goodness is grounded entirely in properties of, or facts about, that thing. For example, the value of human beings, or conscious beings more generally, could be explained entirely by facts about those beings, and no other things.

Extrinsic values, on the other hand, are those whose value is grounded in their own properties and the properties of other things. In this sense, water could be a good candidate for an extrinsic value: its importance seems to rely both on its own physical makeup, and the intrinsically valuable beings who are biologically configured to utterly rely on it. While human beings, or other forms of life, can plausibly have value just because of their properties, it seems doubtful that water would be *valuable* if those beings simply didn't exist.

This seems to do a good job of illuminating the issue at hand. It seems reasonable that those things that rely only on themselves for their value will have it by definition as long as they exist, while those that rely on standing in a certain spatiotemporally contingent relation to such things may not.

So how do we know what is intrinsically valuable? This is the crucial, practical point. Moore, who introduced the recently standard conception of value in terms of intrinsic properties at hand, proposed an isolation technique (1903, 91). According to him, we can review possible candidates by imagining that they exist alone in a universe apart from anything else, and asking if it would still be meaningful to call them good (1903, 91). This typically Moorean approach blends an abstract metaphysical notion with an apparently practical, epistemic method of access. A slightly different gloss is given to the isolation approach by Chisholm (1980) who asks us to instead consider some real thing in our world via mental isolation, and ask whether the attitude we have towards it seems to rely on any other mental considerations.

Now Moore (1903), and many other philosophers in the canon, are happy to identify potentially plausible candidates to survive this sort of test, including wisdom, the good life, justice, or love. Moore specifically settles on interpersonal affection and aesthetic appreciation as the two certain bearers of intrinsic value (1903, 204). However, these sorts of examples only serve to immediately identify the problem at hand: for our definition and associated test to be useful we need to be sure the second guarantees the first.

To illustrate the point further, let's first see if we can be more conservative in our search and end up on obviously safer ground. One thing it seems we can safely say is that all the examples chosen are meant to be relevant to the conscious, rational and living human agent. If they are contingent, it is hard to imagine what else they would be contingent on, apart from the people who live the aforementioned good life; or strive towards and perhaps ultimately possess the wisdom.

I think even this is problematic: while the conclusion is plausible, it also seems to be an open question whether the value of conscious life, or conscious experience, relies in any, perhaps symbiotic way on certain forms of conscious experience. Perhaps, just as I claimed aesthetic experience seems to be contingent on consciousness itself, consciousness relies on aesthetic appreciation among many other things for its own value to really be instantiated.

A popular thought experiment introduced by Robert Nozick (1974, 42-45) asks us to consider a machine that can generate the experience of unlimited, enduring pleasure. He then asks us whether many human beings, when offered the choice between living in the machine or in the real world, would choose the machine. Nozick thinks the answer is no: there is something to our real lives beyond pleasure.

For the sake of this discussion, let's allow for now that lives might be valuable because of their own qualities alone. Nevertheless, as soon as we try and again move away from this starting point, the worry of contingency immediately arises. If it is an open question whether pleasure is really derivative of the good life, for example, then it is an equally open question whether the good life is derivative. In our terms, an open question whether reasons to value the good life have nothing to do with facts about, or properties of, the beings that aspire to it. The problem here is that, instead of asking us what grounding facts are present, we are asked to be sure that other facts *are not* present.

To sum up this first point: intrinsic value as a *definition* about grounding seems like exactly the kind of thing we should want under our approach. The buck-passing approach seems particularly amenable to parsing more or less substantial values via

grounding facts or properties. But an abstract, metaphysical definition is not itself a practical test: we need to understand what epistemic practices are the correct ones to match objects up with that definition. And the isolation approach (whether the stronger Moorean or weaker Chisholmian) is murky in this regard. So my two worries about the intrinsic/extrinsic distinction are:

1. That it may be a misguided definition in the first place: that the view that 'intrinsic' is coextensive with 'cannot vary according to x, y, z' may be based on misleading assumptions.
2. That it is difficult to say whether more than one thing is intrinsically valuable. If this is the case, traditionally atomistic, generalistic views about contingency are actually strengthened, at least when it comes to value. While this doesn't touch on Dancy's holism when applied to rightness, my goal here is to investigate the best possible version of his holism on goodness: and it needs subject matter to operate.

If it seems doubtful whether we can say more than very few, or even any at all things are actually intrinsically valuable, there will be very little work for Dancy and Moore's theories: contextuality is built into the notion of contingent or instrumental values. Pleasure and pain are paradigm examples of obviously good and bad things respectively, but it's easy to find examples of pain perhaps being good and pleasure bad: during medical treatment, or during torture, for example.

When we are told that pain might sometimes be good, or that kindness might be bad, and we expect a strong explanation why, it is puzzling to be told that it is because they are intrinsically valuable or disvaluable, but we don't actually know if they are or

not. Again, if even paradigm examples like pleasure turn out to be contingent, there is no case to answer and we would have limited use for a further, unique theory of contextuality: contextuality is baked in to all kinds of contingent, conditioned, extrinsic or instrumental definitions of value. Even if they are actually specially valuable despite our doubtful approach, it is hard to confidently have the debate: it seems like we need to be pretty certain about where we can apply our theory.

I will now turn to a second, more recent concept that potentially rivals the one at hand for our purposes: that of a default. The notion of a default is one that originates in Dancy's discussion of moral rightness and wrongness, and his engagement with responses to his initial arguments. While the notion of a default reason to *value* is only discussed in a few places in the literature, it is an extension of a more popular debate about default reasons that make an action *right* or *wrong*.

Dancy's original formulation of extreme particularism, which is predicated on the view that *all* moral reasons are holistic, and can count in different ways in different places, or not count at all, was quickly criticised in that it leaves the moral landscape unacceptably 'flat'. In this sense, simply saying that all reasons are holistic does nothing to explain why some things seem to pretty much always count in the same way, while some simply do not.

For example, that an action would involve mass torture and murder seems to pretty much always make it morally wrong, but the fact that an action would involve being mildly rude doesn't. Critics who were willing to accept that reasons could be holistic in both cases still protest that the theory should be more fine-grained in terms of

capturing exactly *how* holistic reasons are. In response, Dancy conceded that we can accept the notion of a 'default', defined as:

...We might think of this notion of a default value as retaining something of what is contained in a pretheoretical notion of intrinsic value, shorn of its invariabilist tendencies... (Dancy, 2003, 638).

Now, what is contained in a pretheoretical notion of intrinsic value, shorn of its invariabilist tendencies, sounds like just the kind of thing we were looking for. Simon Kirchin has reviewed the debate sketched above in his paper *Particularism and Default Valency* (2007) and provided perhaps the most detailed exploration of how the notion should be developed further, specifically on the problematic points we identified above: how we can reliably *know* what things might be defaults so we can take them as the constituents of our discussion.

Kirchin reviews and ultimately discards one possible solution: statistical preponderance. On this view, we can reliably work out what has default valency by looking at many examples of actual moral situations, and noting how often certain features appear. This approach has the immediate advantage of practicality: unlike intrinsic values, which ultimately rely on intuition, statistical preponderance raises the hope that we could *empirically* reach the right answers.

Let's take the example of murder. We could survey historical, archaeological, and anthropological data, as well as current societies, on attitudes towards murder, and putative, objective ethical reasons for and against those attitudes. If we did this, we would certainly find an overwhelming preponderance. To take an evaluative

concept, we could pick some variation of pleasure, happiness, or the good life, and interrogate times and places on the same.

The first issue, according to Kirchin (2007, 28-29) is that this just moves the problem back a layer. Presumably for this to work, it would rely on the hopeful assumption that the moral faculty of the intuitionists exists in the way described: that we really can be reliably sensitive to the presence of goodness at least more often than not. This is because, as realists who have accepted a certain metaphysical take on that realism, for these surveys to be worth anything we would need at least some reason to think people reliably react correctly to evaluative phenomena.

But of course, there is a hopeful response. In our discussion of EBP, we saw that it relies on a much more familiar and uncontroversial reason-faculty, or rationality: something we are used to exercising in the light of reasons in all kinds of areas: whether moral or not, normative or not, justificatory or not. So the question then becomes, can we rely on the aggregated results of many individual reasonings to sketch what has default value? I think this is still difficult, but certainly more promising than Kirchin seems to think. To see if we can improve, let's move to examine more closely the *definition*, or apparent existential status, of the notion of a default.

The first relevant distinction to make is that defaults seem to more accurately capture the phenomena at hand than intrinsic values. The notion of an intrinsic value has long and broad use in philosophy, far outside specific questions about contextual variation: so we are taking something from elsewhere and applying it here because it *seems* promising: surely it is those things that are good because of their own properties alone that are special.

But as philosophers, we can be careful and tactical in advance when we select a definition for a particular purpose: insofar as we seem to understand what part of *the world* we want to refer to, and search for the most accurate term to achieve conceptual capture. If we were starting this debate from a neutral beginning point, it is natural that we would instead begin with the notion of a default. We started off by asking 'what shouldn't be able to change in value?'. This is baked in to concept of a default, by definition. Insofar as we wish to use the most accurate and fitting term to refer to a particular thing, a default is surely a default because it ought not to regularly vary. It seems impossible that it would be *less* useful as a definition than that of an intrinsic, and very possible that it is *more* fitting.

The second relevant distinction is that between binary notions and probability-based notions. The intrinsic/extrinsic distinction is a binary one: something is either good or it is not. This creates a perhaps impossibly high standard of proof, as per our discussion above: ultimately one would have to claim to have surveyed *the entire world* to be able to rule with certainty on precisely what grounds what. Defaults, on the other hand, have a lower barrier to entry: a default only has to be the relevant way *usually* to qualify, and one doesn't necessarily need to be certain about what grounds what: defaults give probabilistic answers to a probabilistic question.

The third point adds an extra barrier of proof to the first: to make a claim about intrinsic value we have to include a claim about the *grounds*, whereas with defaults we only have to mention reasons and attitudes. This is important because attitudes and the reasons that inspired them are much easier to survey. If we carried out a survey of the kind Kirchin hints at, asking people's attitudes towards certain things

would be easier to accept because of the point above: they would potentially have had to survey the entire world, instead of just their own considered attitudes.

Kirchin tells us that a better approach is to take the definition absolutely directly: defaults are those features that reliably only seem to change according to enabling features, so by cataloguing instances where we are sure there are no such secondary considerations, we can build up a default map. But this also suffers according to the issue about cataloguing grounds. If we can't be utterly certain no other grounds are present, how can we do any better regarding non-grounding enabling features? In both cases these things refer to ordinary facts. If grounding is always potentially global, then enabling and disabling are always potentially global.

Ultimately, I think statistical preponderance turns out to be the far more workable and practical option. Remember the point is not that statistical preponderance is wrong: in fact it is guaranteed. By definition, a default consideration has to present in that form in the majority of places. Kirchin's point is an epistemic one: he questions whether we grasp those reasons correctly often enough to accurately capture the probabilities.

In conclusion, the notion of a default is both more fitting to that pre-theoretical core Dancy refers to, and easier to accept and put into practice on epistemic grounds. If I want to know what things entirely ground their own value, and carry it about with them, I can ask what has intrinsic value. If I want to know what are the fitting things to orient myself towards, what I should fundamentally value, I can ask what has final value. If I want to know what things are more resistant to contextual change than others, I can ask what has default value. And in doing so, I can avoid the perhaps

impossibly high epistemic barriers to having certain knowledge of intrinsic values: all I need to know is whether the overwhelming mass of evaluations have clearly pushed in a certain way. And this seems to be how a discussion like this begins in the first place: before we ever approach philosophy, it seems certain to us that murder is a default wrong, or pleasure a default good. I will now move to consider our third question: how such a default might still be contextualised.

Chapter Five : Organic Unity and

Value Holism

I have now advocated for what I say are the most attractive extant theoretical answers to our first two questions: a buck-passing explanation of what it means to be good, and a 'defaults' explanation of what it means to be good in a special way. I will now turn to our third and final question: the most attractive theoretical explanations of how such a default value can vary: e.g. in valence. There are two such theories: Moore's doctrine of organic unities (1903, 30-33), and Dancy's doctrine of value holism (2003, 629-641, 2004b, 176-189, 2008, 97-99). The two positions have variously been put into agreement (e.g. Mckeever & Ridge, 2013) and disagreement (e.g. Brown, 2007) in the literature.

I will first say something again about what I mean by values changing, or contextuality. On reflection, there are three things this can mean. Firstly, it can refer

to whether something is the subject of evaluation at all: whether it is even potentially or actually good or bad, or simply nothing. Zimmerman (2015, 10) calls this evaluative adequacy or inadequacy: whether something meets the minimum standards to be evaluated; Moore refers to the same notion in terms of there being a class of things that are evaluatively 'indifferent' (1903, 27). This point becomes interesting if we ask whether such a thing can in some contexts become valuable; or contribute to the value of something else. The second point is one about valence: whether the same thing can be good in one place yet bad in another, or bad in one place yet contribute to the goodness of something else in another. The third point is one about amount. Insofar as we might say one thing can be more valuable than another, or that something can become more valuable, we might ask whether the same thing can gain value, or contribute to a gain in value, from place to place.

There are two strands here: one about single objects and one about arrangements. The questions about single objects are those about one thing being evaluable at all, or if so, valenced. The questions about arrangements are those about things that are indifferent, good, or bad contributing to wholes with different evaluable statuses or valences. In both cases, amounts are important: the amount of value any particular thing carries about with it, or the amount it can contribute to a whole, which also has a certain amount.

Dancy's view focuses on the conditions under which a single, distinct object might vary in value, while Moore's is of the second type: it tells us all about how multiple values or disvalues can interact in contextually organic ways. I will begin by simply outlining both views in their original forms, without referring to my earlier

complaints, before turning to the two main problems identified in the literature. This is because my argument for rejecting these objections relies on adjusting the views in several ways, and this needs to be drawn out in order.

The first of the two extant and distinct explanations of evaluative contextuality is Moore's doctrine of organic unities. This view holds that certain arrangements of parts, which have distinct evaluative identities, have a master identity that can be more or less valuable than the sum of the parts.

This is a view that, in common with others discussed, has intuitive pre-theoretical appeal and can be initially motivated with ordinary non-moral examples. In our everyday lives, that something can be 'greater than the sum of its parts' is a truism, according to the definition above. A car has more value than the individual nuts and bolts, a house more than the bricks, a song more than the individual notes, and so on. Equally, the notion that a whole can be worth *less* than the sum of its parts is easy to appreciate: consider popular debates over adding pineapple to pizza. Dancy (2004b, 181) uses an example drawn from aesthetic appreciation of a dress and a valuable diamond: which are separately attractive but clash horribly together.

The second horn of Moore's position is that a part that has positive value can contribute to a negative whole, or vice-versa; and that parts that simply have no value at all (are evaluatively inadequate) can contribute to parts that are evaluatively valenced, and this is similarly easy to illustrate. Two simple examples of this sort of contingency are criminal punishment, and surgery. In these cases, suffering, which is presumably a 'default' bad thing, contributes to an apparently positive outcome. An example of the reverse is where pleasure seems to contribute to a negative whole:

the pleasure of a torturer, say, or more simply, the *schadenfreude* we sometimes feel at another's misfortune.

Theoretically, Moore's achievement is to extend these considerations to weightier, distinctly *ethical* questions of the same form, and to attempt to give a firm analytic foundation to the notion of organic unities. For Moore, questions like punishing the guilty are of the same form of those above in that they correctly identify something that needs explanation, but differ in that the required explanation is of a different form. This is because, again for Moore, ethical goodness and badness robustly exist.

So, the doctrine of organic unities tells us that 'wholes' can be greater or less than the sum of their parts: Moore uses the example of the human creature, whose whole has more value than the sum of its arms, legs, organs etc. It also tells us that wholes can have an evaluative valence different than the valence of some of the parts taken individually. Moore uses the example of just punishment: while pain and violence are both bad, the net state of affairs when punishing a serial killer may be a good one because of the value of lots of other things. In either case, something that has no value at all of its own can contribute to the goodness or badness of something.

This view is holistic in the sense that it asks us to focus on the whole of the apparently relevant objects in any particular case. So Moore's view seems to cover all three of the dimensions of contextuality I pointed out: whether something is valuable at all, how valuable it is, and what valence it has. Dancy's view is holism about value, and is expressed as 'anything that is good in one place, can be bad, or nothing at all, in another place'. This is explained via another claim: that features that are not values or grounds for values can influence the evaluative status or valence of things that

are. This view is holistic in a broader sense than Moore's: it asks us to consider not just the features of objects of value in a situation, but all of the features of any kind at all, potentially including the entire world. As Dancy phrases it, supervenience is always potentially global, whereas for Moore it is local (2004b, 178).

So, we have two initially compelling views about contextual evaluation that overlap in some ways but not in others. Why are these views held to be incompatible? There are two strands in the literature in this regard: additivism and variabilism (Brown, 2007). Additivism (also referred to as combinatorialism (Berker, 2007)) is the view that values and disvalues come in set amounts that can be conceptualised, or modelled, numerically. On this view, a certain instance of pleasure might have a 'value' of +4, while a certain instance of pain might have a 'disvalue' of -4, the two together forming an equilibrium. This means that, at least in theory, considering the final value of any collection of goods and bads is a matter of evaluative arithmetic.

Moore rejects combinatorialism and holds that values only have metaphorical amounts that cannot be expressed as described: weighing values is instead stubbornly a matter of sensitive appreciation. Dancy and others defend combinatorialism about values. Specifically, Dancy identifies two strands of Moore's argument about wholes: that a part can contribute *more* or *less* value than it has, respectively. He identifies the claim about only contributing part of a parts value as unproblematic. Using the example of a person, Dancy points out that a person can contribute goodness to the world as a friend, a parent, or a colleague, and that they are surely greater than any particular part in this regard. Instead, he problematises the positive claim.

Invariablism is the claim that any particular thing, when taken alone, cannot change in value. Notice again that Moore's theory only discusses arrangements of good or bad things: he says nothing about *the same thing* alone being able to change. Dancy, however, advocates for precisely that: his holistic theory tells us that secondary features that are not part of objects of value can have such an effect.

So above I identified how each theory brings something unique to the table, and that it would be good to be able to endorse both, precisely because one deals with wholes and one with individual objects. The combined theory could be a complete and internally consistent 'holistic' take on value. But to be able to do that, we have to deal with combinatorialism and invariablism. I will argue that we should reject both, and that we can do so precisely because we also endorsed EBP.

Combinatorialism is the simpler of the two to deal with, so I shall take that one first. Dancy endorses combinatorialism precisely because of his view that value is a *necessarily* robust and distinct property, which things have in certain amounts. Remember that while Dancy endorses a reason-based evaluative epistemology, he holds that values and reasons are coextensive: when we get it right, the strength of the reasons we have will be in accordance with how valuable the thing really is; how much of the value property it has.

This is why he claims that on Moore's view, objects can contribute value they 'have not got' (2003, 3) and that this violates the necessary connection between values and reasons (2003, 4). On his view, the extra value that appears in the situation apparently appears out of thin air, since it cannot have come from the value the thing has.

So, Dancy essentially rejects the positive horn of Moore's argument for the reason he is happy to endorse the second: it respects that necessary relation. For Dancy, it makes complete sense that an object only contributing part of its value will engender weaker reasons to endorse it: the properties are still coextensive. But in the other sort of case, telling us that arrangement *x* is more valuable than individual objects *y* and *z* surely ought to mean that we *have more reason* to value it. But if values and reasons spring up coextensively from the same ground, then there will be a concomitant increase in the value properties of the parts in line with the whole.

This objection is surely plausible. Luckily, we can tackle it in a straightforward way: via reference to our earlier defence of EBP, and Dancy's own work. On EBP, there is no value property that has the possibility of disagreeing with reasons: since to be valuable *just is* to have reason-giving properties, distinct co-extensiveness isn't present in the ontology. The question then becomes one about whether *reasons* seem to have weights; and Dancy argues that they do not: indeed, he has been criticised for advancing this point so consistently (Berker, 2007). So combinatorialism becomes unproblematic on two levels, at least as far as reconciling it with Dancy is concerned. Not only can we hold that reasons don't have quantifiable strengths in the first place, but we can also hold that they can vary in that strength even when the essential ground remains unchanged, presumably because of other facts that can influence reasons as enabling factors being present in the situation, which for Dancy can be construed incredibly broadly; remember again that for Dancy supervenience is always potentially global, and any facts or properties in the universe can play this role.

This leads us neatly into the second problem: the notion that the *valence* of the reasons pertaining to an object, or the evaluative adequacy or inadequacy of that object, can vary even when that essential ground remains unchanged. The debate has now reversed: Moore rejects the notion that a single object can change in value as strongly as Dancy rejects the principle of organic unity about arrangements. This is a claim that Dancy defends in relation to both goodness and rightness, and in the case of goodness, which is our focus, his position is completely preserved despite endorsing EBP.

As we have seen previously, the way Dancy motivates this is via what he calls secondary features: relevant facts that are not reasons or the grounds of reasons but can weigh on the valence or existence of reasons (Dancy, 2004b, 38-52). Just as he extends his reason holism to his notion of robust value, so too does he extend the notion of a default: elsewhere, he claims that even the most valuable things can be affected by such contextual factors (Dancy, 2003, 632). However, we have no need to explore this second claim directly: by endorsing EBP the first claim is completely sufficient.

Now, Moore's work precedes Dancy's by a full century and therefore he does not comment explicitly on the plausibility of Dancy's specific view about contextual factors. All we can do here is consider why Moore endorses invariablism on his own terms, and this is a combination of two views we have already examined and problematised: the Moorean property of goodness, and a focus on *intrinsic* value. On the first point, accepting EBP means rejecting that property, and rejecting any view that predicates invariablism on that property by definition. Similarly, I problematised

how we can conceptualise intrinsic value and how certain we can be about what has it.

This puts us in a position where, at least for our own purposes, defending a contradiction between Dancy and Moore cannot succeed simply by reiterating Moore's positive claims. Instead, we can only make a negative claim followed by a modest conclusion: that Dancy's view on secondary factors is wrong, and therefore the variability/invariablism debate is still open.

So ultimately bringing Moore into conversation with Dancy involves endorsing his version of holism about default reasons: the view that any reason can on principle vary because of contextual factors, and that in the case of a default, these contextual factors must be extraordinary. While I will not directly survey all the objections to the core holistic thesis here, suffice to say that if it is independently plausible (and I think it is) then there is no problem with rejecting invariablism. My goal here has been to sketch a plausible picture of how these two independently interesting and plausible axiological views can be reconciled, and in my focus on EBP, I hope I have offered an extra, interesting way of reconciling the views, alongside those philosophers who have also advocated for that conclusion, and against those who hold that the views are necessarily mutually exclusive.

Conclusion

In this thesis I have shown that the contextual views of Moore and Dancy are not, as some commentators suggest, incompatible or antagonistic, but instead channel the same practical spirit into explaining evaluative phenomena: relating to single objects and collections respectively. I argued that we can do so precisely because we accept the buck-passing view about value, and in making this original contribution to a developing literature I feel I can channel a perhaps lighthearted sentiment of Mckeever & Ridge (2013, 280); That the Moorean and Dancyean accounts together seem to form an organic unity, and if this is true, and my account is also true, then the views of Moore, Dancy and Scanlon together form an even more attractive organic unity.

I have attempted to ground my conversation in ordinary, accessible concerns precisely because this is what is admirable about the Dancyean project, which in this sense channels the spirit, if not the letter, of the Moorean project that began the recent tradition it aims to continue. Both Moore and Dancy begin with a practical, common-sense interest in ethics and evaluation that is translated 'upstream' into formal meta-discussion that aims to not lose sight of the concerns that drove it, and succeeds when it can be coherently brought back down to that practical level.

Beginning with the work of J.L. Mackie, I attempted to preserve what is interesting and plausible about non-naturalism while tempering some of the more extravagant forms robust views can take. It is precisely because my focus on non-naturalism is inspired by pre-theoretical, 'common-sense' concerns that weigh on real issues that I have sought to review the views at hand in the way I have, by seeking the simplest, most broadly applicable, parsimonious and acceptable explanations.

Similarly, my defence of a more practical, yet equally interesting notion of default value over intrinsic value is done with a keen focus on practical over theoretical concerns. I have tried to demonstrate that the concept of a default value fits well with what we really want to talk about, and actually allows us to have a meaningful conversation: discussions of intrinsic value, particularly on what epistemic practices allow us to apprehend that value, are at their best when they keep this in mind, and at their worst where they wander into abstractions of dubious applicability, at least when discussing contextuality in ethics or evaluation.

In conclusion, accepting these views expands the horizons of Dancy's work insofar as on the buck-passing view, it specifically allows equally adequate explanatory power with less distinct metaphysical categories, while more broadly contributing to the promotion of not just an ethical, but normative explanation. On the Moorean view, reconciling Dancy and Moore specifically gives Dancy's holism more explanatory power and potential in that Moore's insights capture phenomena Dancy's do not, while the combined package is a strong position in debates about ethical, evaluative or normative contextuality.

This thesis is intended to form part of a longer research arc of larger scope. Particularism is a relatively recent and controversial interjection into traditional debates on *right action*, and this is naturally what Dancy devotes the most pages to, as do his critics. Particularism has variously been criticised as problematic regarding explanations of interpersonal moral justification, and of moral education, for example. The problem here is that full practical explanations of either of these things will ultimately be predicated on a fully-fledged *ethical theory*.

And particularism is not an ethical theory at all. It is a relatively small cluster of meta-level claims, that are officially theory-neutral. While it is an open question whether normative ethics will expand in the future, with further distinct sorts of views available, in recent history philosophers have coalesced mostly around two key and wide-ranging approaches: consequentialism and deontology. For the sake of what I am about to say, let's assume that one of these two views will turn out to be correct.

The distinctions between the two views, and whether either, neither or both are held to be potentially plausible, are necessarily derived from the meta-level. Both approaches require a full meta-theory of axiology and deonticity to be fulfilled: with each of them focusing particularly on one of these conceptual areas. So if particularism is to be similarly drawn down, and some distinctly particularist theory derived at the first-order level, it ought to have robust meta-theories about both evaluation and action. I believe that particularists have done better on the second point than the first, and that particularist axiology therefore deserves particular attention. In this essay I have tried to make some useful points on what the evaluative part of the particularist meta-philosophy could look like, in a way that is meant to naturally and strongly accord with their better-known and more finely developed theories of action.

Bibliography

Allen, J. *G.E. Moore and the Principle of Organic Unity*. In *The Journal of Value Inquiry*, No. 37, Vol. 1, 2003. Pp. 329-339.

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.

Baker, A. *Quantitative Parsimony and Explanatory Power*. In *British Journal for the Philosophy of Science*, Vol. 54, No. 2, 2003. Pp. 245-259.

Brown, C. *Two Kinds of Holism About Values*. In *the Philosophical Quarterly*, Vol. 57, No. 228, 2007. Pp. 456-463.

Bedke, M. S. *Intuitive Non-Naturalism Meets Cosmic Coincidence*. In *Pacific Philosophical Quarterly*, Vol. 90, No. 2, 2009. Pp. 188-209.

- *Might All Normativity Be Queer?* In *Australasian Journal of Philosophy*, Vol. 88, No. 1, 2010. Pp. 41-58.

Beebe, J. *The Empirical Study of Folk Metaethics*. In *Etyka*, Vol. 15, No. 1, 2015. Pp. 11-28.

Bergqvist, A. Cowan, R. *Evaluative Perception*. [Oxford: Oxford University Press, 2018].

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

- *The Unity of Grounding*. In *Mind*, vol. 127, No. 507, 2018. Pp. 729-777.

Branmark, J. *Good-Making and Organic Unity*. In *Philosophical Studies*, Vol. 174, No. 1, 2017. Pp. 1499-1516.

Brink, D. O. *Moral Realism and the Sceptical Arguments from Disagreement and Queerness*. In *Australasian Journal of Philosophy*, Vol. 62, No. 2, 1984. Pp. 111-125.

Broome, J. *Reason Fundamentalism and What Is Wrong with It*. In *the Oxford Handbook of Reasons and Normativity*. Ed. Star, D. Pp. 297-318. [Oxford: Oxford University Press, 2018].

Bykvist, K. *No Good Fit: Why the Fitting Attitude Analysis of Value Fails*. In *Mind*, Vol. 118, No. 469, 2009. Pp. 1-30.

Chisholm R. *Defining Intrinsic Value*. In *Analysis*, Vol. 41, No. 2, 1981. Pp. 99-100.

Cook, R. *An Augmented Buck-Passing Account of Reasons and Value: Scanlon and Crisp on What Stops the Buck*. In *Utilitas*, Vol. 20, No. 4, 2008. Pp. 490-507.

Crisp, R. *Goodness and Reasons: Accentuating the Negative*. In *Mind*, Vol. 117, No. 466, 2008. Pp. 257-265.

- *Value... And What Follows by Joel Kupperman*. In *Philosophy*, Vol. 75, No. 3, 2000. Pp. 452-462.

Dancy, J. *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*. [Clarendon: Oxford, 2004a]

- *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.

- *Should We Pass the Buck?* In *Royal Institute of Philosophy Supplements*, Vol. 47, No. 1, 2000. Pp. 159-173.

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

Danielson, S. and Olson, J. *Brentano and the Buck-Passers*. In *Mind*, Vol. 116, No. 463, 2007. Pp. 511-522.

Darwall, S. *Philosophical Ethics*. [Colorado: Westview Press, 1998].

Feldman, F. *The Open Question Argument: What it isn't, and What it is*. In *Philosophical Issues*, Vol. 15, No. 1, 2005. Pp. 22-43.

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.

Fletcher, G. *Brown and Moore's Value Invariabilism vs Dancy's Variabilism*. In *The Philosophical Quarterly*, Vol. 60, no. 238, 2010. Pp. 162-168.

Frankena, W. K. *The Naturalistic Fallacy*. In *Mind*, Vol. 48, No. 192, 1939. Pp. 464-477.

Hare, R. M. *The Language of Morals*. [Oxford: Oxford University Press, 1964].

Heathwood, C. *Fitting Attitudes and Welfare*. In *Oxford Studies in Metaethics*, Vol. 3, No. 1, 2008. Pp. 47-73.

Hooker, B. and Stratton-Lake, P. *Scanlon versus Moore on Goodness*. In *Metaethics After Moore*. Ed. Horgan, T. and Timmons, M. [Oxford: Oxford University Press, 2006]. Pp. 149-168.

Hurka, T. *Two Kinds of Organic Unity*. In *The Journal of Ethics*, Vol. 2, No. 4, 1998. Pp. 299-320.

- *Asymmetries in Value*. In *Nous*, No. 44, Vol. 2, 2010. Pp. 199-223.

Jackson, Frank. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. [New York: Oxford University Press, 2000].

Kim, J. *Moral and Natural Kinds: What's the Difference For a Naturalist?* In *Philosophical Issues*, Vol. 8, No. 1, 1997. Pp. 293-301.

- *Supervenience as a Philosophical Concept*. In *Metaphilosophy*, Vol. 21, No. 1-2, 1990. Pp. 1-27.

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

Korsgaard, C. *Two Distinctions in Goodness*. In *The Philosophical Review*, vol. 92, no. 2, 1983. Pp. 169-195.

Lang, G. *The Right Kind of Solution to the Wrong Kind of Reason Problem*. In *Mind*, Vol. 20, No. 4, 2008. Pp. 472-489.

Lemos, N. *A Defence of Organic Unities*. In *The Journal of Ethics*, Vol. 19, No. 2, 2015. Pp. 125-141.

Lillehammer, H. *Companions in Guilt: Entailment, Analogy, and Absorption*. In *Companions in Guilt Arguments in Metaethics*, Ed. Cowie, C. and Rowland, R. [Routledge: London, 2019].

Little, M. *In Defence of Non-Deontic Reasons*. In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, Ed. Bakhurst, D. Hooker, B. and Little, M. Pp. 112-136. [Oxford: Oxford University Press, 2013].

Mackie, J. L. *Ethics: Inventing Right and Wrong*. [London: Penguin, 1977]. Pp. 38-42.

McKeever, S. and Ridge, M. *Why Holists Should Love Organic Unities*. In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, Ed. Bakhurst, D. Hooker, B. and Little, M. Pp. 265-285. [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.

Metz, E. *How to be a Deontic Buck-Passer*. In *Philosophical Studies*, No. 77, 2020. Pp. 3193-3211.

Mintz-Woo, K. *On Parfit's Ontology*. In *Canadian Journal of Philosophy*, Vol. 48, No. 5, 2018. Pp. 707-725.

Moore, G. *Ethics, 1912*. [New York: Oxford University Press, 1965].

- *Principia Ethica, 1903*. [Ebook: Digireads, 2012].

Morton, J. and Sampson, E. *Parsimony and the Argument from Queerness*. In *Res Philosophica*, Vol. 91, No. 4, 2014. Pp. 609-627.

Nozick, R. *Anarchy, State, and Utopia*. [Oxford: Blackwell, 1974].

Olson, J. *Buck-Passing and the Wrong Kind of Reasons*. In *The Philosophical Quarterly*, Vol. 54, No. 215, 2004. Pp. 295-300.

- *G.E. Moore on Goodness and Reasons*. In *Australasian Journal of Philosophy*, Vol. 84, No. 4, 2006. Pp. 525-534.
- *The Metaphysics of Reasons*. In *In the Oxford Handbook of Reasons and Normativity*. Ed. Star, D. [Oxford: Oxford University Press, 2018].

Parfit, D. *On What Matters*. [Oxford: Oxford University Press, 2011].

Piller, C. *Content-Related and Attitude-Related Reasons for Preferences*. In *Royal Institute of Philosophy Supplement*, Vol. 59, No. 1, 2006. pp.155-182.

Potrc, M. *Particularism and Resultance*. In *Acta Analytica*, Vol. 19, No. 1, 2004. Pp. 163-187.

Rabinowicz, W. & Rønnow-Rasmussen, T. *The Strike of the Demon: On Fitting Pro-attitudes and Value*. In *Ethics*, Vol. 114, No. 3, 2004. Pp. 391-423.

Reisner, A. *Not Fittingness, not reasons, and no Value: Against the 'First' Views*, 2021. (unpublished).

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

Ross, W. D. *The Right and the Good, 1930*. Ed. By Stratton-Lake, P. [Oxford: Oxford University Press, 2002].

Rowland, R. *Reasons or Fittingness First?* In *Ethics*, Vol. 128, No. 1, 2017. Pp. 212-229.

Scanlon, T. *What We Owe to Each Other*. [Massachusetts: Harvard University Press, 1998].

Schroeder, M. *The Fundamental Reason for Reasons Fundamentalism*. In *Philosophical Studies*, 2020 (forthcoming).

Sendlak, M. *On Quantitative and Qualitative Parsimony*. In *Metaphilosophy*, Vol. 49, No. 1-2, 2018. Pp. 153-166.

Shafer-Landau, R. *Moral Realism: A Defence*. [Oxford: Clarendon, 2003].

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

Skorupski, J. *Buck-Passing about Goodness*. In *Hommage a Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowitz*. Eds. Ronnow-Rasmussen, T. Petersson, B. Josefsson, J. Egonsson, D. 2007.

Smith, M. *The Ideal of Orthonomous Action, or the How and Why of Buck-Passing*. In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, Ed. Bakhurst, D. Hooker, B. and Little, M. Pp. 50-75. [Oxford: Oxford University Press, 2013].

Stratton-Lake, P. *Dancy on Buck-Passing*. In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, Ed. Bakhurst, D. Hooker, B. and Little, M. Pp. 76-96. [Oxford: Oxford University Press, 2013].

- *How to Deal with Evil Demons: Comment on Rabinowicz and Ronnow-Rasmussen*. In *Ethics*, Vol. 115, No. 4, 2005. Pp. 788-798.

Suikkanen, J. *Buck-Passing Accounts of Value*. In *Philosophy Compass*, Vol. 4, No. 5, 2009. Pp. 768-779.

- *In Defence of the Buck-Passing Account*. In *Ethical Theory and Moral Practice*, Vol. 7, No. 5, 2005. Pp. 513-535.

Toppinen, T. *Is Irreducible Normativity Impossibly Queer?* In *Journal of Moral Philosophy*, Vol. 13, No. 4, 2016. Pp. 437-460.

Turner, J. *Curbing Enthusiasm About Grounding*. In *Philosophical Perspectives*, Vol. 30, No. 1, 2016. Pp. 366-396.

Vayrynen, P. *Resisting the Buck-Passing Account of Value*. In *Oxford Studies in Metaethics*, Vol. 1, No. 1, 2006. Pp. 295-324.

Wodak, D. *Who's on First?* In *Oxford Studies in Metaethics*, Vol. 15, No. 1, 2020. Pp. 49-71.

Zimmerman, M. *On the Nature, Existence and Significance of Organic Unities*. In *Journal of Ethics and Social Philosophy*, Vol. 8, No. 3, 2015. Pp. 1-25.

