

Robustness in Moral Reality

Norms, Necessity, and Robust Realism

Stephen Ingram

Thesis submitted for the degree of Doctor of Philosophy
(Philosophy)

March 2016

Department of Philosophy
University of Sheffield

Abstract

This thesis examines the metaethical theory known as 'Robust Realism.' According to defenders of this view, there exist irreducible, non-natural, mind-independent, and categorically authoritative moral properties. The central aim of this thesis is to identify the best way of understanding and motivating these claims. In other words, I intend to develop a compelling metaphysics for Robust Realism. I don't plan to show that Robust Realism is true, but I do plan to identify the best formulation of it. I will thereby put us in a better place to assess its viability against rival views of moral reality.

The robustly realistic theory that I will develop is built around the idea that there are necessary moral norms. In other words, norms that have authority in every possible world. I show how positing such norms enables the Robust Realist to defuse two influential ways of objecting to their claim that moral properties are irreducible. I provide an account of necessary moral norms as fundamental entities with a modal jurisdiction that, unlike the many non-fundamental moral norms, is not limited by any contingent presupposition. I show that the mind-independence of moral properties and norms takes us some way toward the elimination of those contingent limiting factors. I use this account to clarify the categoricity of moral direction, and in turn show how the categoricity of moral direction can be used in defence of an interestingly non-naturalist view of moral reality.

I thereby give a compelling metaphysics for Robust Realism, but I do not thereby show that this theory is true. I thus consider prominent ways of arguing about moral reality, to assess whether we can decide the matter one way or another. Unfortunately, however, I show that debate about moral reality often results in a persistent stalemate. I diagnose this by appealing to deep differences in 'temperament' and 'existential need.' I thus conclude that we might need to limit the ambitions of metaethical inquiry.

Acknowledgements

I pretty much know who to thank, but I'm much less confident about how to thank them. I owe many people a great deal, however, so I will do my best.

Jimmy Lenman and Miranda Fricker have both been wonderfully encouraging supervisors, and over the last few years they have read most of my written work – an unenviable undertaking. The content of my thesis has been improved considerably by their incisive questions, suggestions, and ideas. I am grateful to Miranda for providing insightful new ways of thinking about difficult questions, and for helping me to sharpen my view of certain methodological issues. I am particularly grateful to Jimmy – my primary supervisor – for his time, his intellectual charity, his perceptive observations, and his general support in helping me to produce an actual thesis.

Suki Finn and Graham Bex Priestley gave comments on an entire draft of my thesis, despite having no legal obligation to do so. I am grateful to both of them for their comments, which have been edifying. I am also grateful to Neil Williams and Ian Kidd, both of whom gave very helpful comments on a draft of Chapter 6. They also suggested interesting things to read that I might have otherwise missed. Aaron Wilson read a paper that failed to make it even into the earliest of early drafts, but that was written at a transitional point in the course of my PhD. For reading that paper with rather more charity than it most likely merited, I'm grateful to him.

There are many friends who have, in various ways, made Sheffield a marvellous place to be. The following excludes those mentioned above, and is in reverse alphabetical order: Steve Wright, Jack Wadham, Abbey Southall, Poppy Shelton, Joe Saunders, Nick Rebol, Phil Rau, Angie Pepper, Ashley Pennington, Jonathan Parry, Siobhan Moriarty, Robbie Morgan, Natasha McKeever, Neri Marsili, James Lewis, Damiano La Manna, Simon Kittle, Armin Khameh, Nicola Kemp, Tom K, Katharine Jenkins, Rich Healey, Toni Harrison, Trystan Goetze, Denise Fox, Carl Fox, Ahmad Fattah, Anton Eriksson, Alex Duval, Ryan Doran, Charlie Crerar, May Connolly, Adriana Clavel, Shirley Carter, Pete Caven, Lewis Brooks, Stephen Bolton, Josh Black, Jess Begon, Simon Barker, Al Baker, and Charlotte Alderwick. It is also a joy to see friends beyond Sheffield, like Holly Savage, Paul MacDonald, Michael Harris, Conrad Brown and Tom Becker.

I'd also like to thank all of the staff in the philosophy department, for helping to make Sheffield a marvellous place to be a student. The administrative staff are endlessly helpful and, impressively in some cases, patient. The academic staff are invariably generous with their time, their advice, and their encouragement. They have provided a model of collegial conduct from which I have learned a great deal, so I am grateful to all of them.

In the last three years I have had the opportunity to speak at a number of seminars, conferences, and workshops. My views on a variety of issues related to this thesis have been clarified significantly as a result of the constructive feedback received on those occasions. So, in addition to participants in the graduate and summer seminars at Sheffield, I would like to thank audiences at Manchester, Hull, Durham, Southampton, and Leeds for their willingness to engage charitably with often undercooked ideas. I'd also like to thank to the Arts and Humanities Research Council for funding my project. Their financial support has made life as a graduate student a great deal less stressful than it would otherwise have been, and has enabled me to devote most of my time and attention to the thinking, reading, talking, and writing that pretty much constitute a philosophy PhD.

I would not have reached this point without my parents, Maggie and Dave, who passed on a firm belief in the importance of education, along with an ethic of hard work that has been vital in the last few years. Their support has been diverse in form, but unified in constancy. I am grateful to both of them. The same goes for Lizzy Kirkham. I made several attempts to put my thoughts about her into words, but it's a rather tricky thing to do. Sincere expressions of love and admiration aren't really my style. I wanted to write something clever or funny, but I couldn't think of anything that didn't seem too superficial or glib. I tried to write something sincere and heartfelt, but it just felt solemn or earnest. I tried to quote something wise and deep, but nothing felt apt, and I'm not quite literate enough to execute such a move without affectation. I therefore settle for this awkward (but I will assume endearing) review of my various failed efforts.

Contents

1	Starting Points	1
1.1	Robust Realism	2
1.2	Metaphysics	9
1.3	Metaphysics and Robust Realism	12
1.4	Necessities	16
1.5	Spoilers	18
2	The Possibility of Robust Realism	21
2.1	Two Points of View	22
2.2	Anti-Archimedean Challenges	24
2.3	Quasi-Realism	26
2.4	Foundations of Quasi-Realism	29
2.5	Relaxed Realism	32
2.6	Compatibility	36
2.7	Against Compatibility	40
2.8	Moral and Metaphysical Argument	45
2.9	Conclusion	48
3	The Supervenience of Moral Properties	49
3.1	Non-Reductivism and Descriptivism	49
3.2	Moral Supervenience	54
3.3	Reductive Supervenience Objections	56
3.4	Necessity and Redundancy	61
3.5	Explanatory Supervenience Objections	64
3.6	Necessity and Norms	66
3.7	Norms and Explanations	70
3.8	Necessary Moral Truths	73
3.9	Conclusion	78
4	The Necessity in Robust Realism	79
4.1	Moral Fixed Points	79
4.2	Conceptual Deficiencies	83
4.3	Against Moral Fixed Points	85
4.4	Moral Norms	88
4.5	Limiting Factors	90
4.6	A Case Study: Mind-Independence	94
4.7	More on Mind-Independence	98
4.8	Necessary Norms in Ordinary Life	101
4.9	Conclusion	104

5	The Authority of Morality	105
5.1	Images of Authority	106
5.2	Categoricity and Practical Necessity	109
5.3	Categoricity and Necessitation of Blameworthiness	111
5.4	Non-Naturalism and Metaphysics	116
5.5	Non-Naturalism and Methodology	119
5.6	Beyond Naturalism: Part I	122
5.7	Beyond Naturalism: Part II	126
5.8	Arguing from Queerness	131
5.9	Conclusion	134
6	Arguing about Moral Reality	135
6.1	The Substantive Strategy	135
6.2	Stalemate	140
6.3	Temperament	142
6.4	Existential Needs	145
6.5	The Debunking Strategy	149
6.6	Good Reason	152
6.7	The Temperament in Debunking	154
6.8	The Limits of Metaethics	157
6.9	Conclusion	160
7	Conclusion	161
	References	163

1 Starting Points

You have to start somewhere, and there are a number of places at which one might begin an investigation into the nature of morality. Some begin by asking about the content of our moral terms, like ‘right’ and ‘good’ and ‘ought.’ Others try to make sense of the motivational role played by moral judgement, or they seek to improve our knowledge of human nature and moral psychology. These are respectable starting points for ‘metaethical’ inquiry, that is, inquiry into the nature and status of ethics. Semantic and psychological issues are interesting and important. But they are not my starting point. I begin with metaphysics. I want to make sense of moral reality.

More specifically, I want to answer the following question: what is the best way to understand the metaphysical commitments incurred by robustly realistic metaethical theories? In answering this question I will develop a set of metaphysical positions that should be accepted by those sympathetic to a view known as ‘Robust Realism.’ Developing this set of positions is the aim that guides subsequent chapters. But what is Robust Realism? I will answer this question thoroughly in §1.1, but here is a first approximation: Robust Realism is the claim that there is an irreducible, non-natural, mind-independent, categorically authoritative moral reality. Although it was long considered dead and buried, Robust Realism has in recent years been exhumed and revived by a number of philosophers.¹ It is now alive and kicking, and I intend to advance our understanding of it.

I don’t intend to show that Robust Realism is true. For reasons that will eventually become clear, I’m simply not that ambitious. My aim is quite modest, for I just want to show how a robust moral metaphysics should be developed. However, despite its relative modesty, this aim is important. There are many ways to view the metaphysical commitments of Robust Realism, and we won’t be in a position to properly compare this theory with its rivals until we have found the most convincing version of it.

¹ Contemporary defenders include Shafer-Landau (2003), Huemer (2005), Oddie (2005), Wedgwood (2007), Cuneo (2007a), FitzPatrick (2008; 2011; 2014), Tännsjö (2010), Enoch (2007; 2011), and Wielenberg (2014). Not all of these philosophers use the term ‘robust,’ though some do, and all are Robust Realists in the sense given in §1.1. Also, some discuss normativity in general rather than morality in particular. I will concentrate mainly on morality, as that’s what I’m interested in and you have to start somewhere.

So, in investigating the nature of morality, my starting points are metaphysics and, more specifically, the metaphysics of Robust Realism. One might doubt that this is the best way to begin a metaethical inquiry. Many are suspicious of metaphysics in general, and even those who aren't might nevertheless regard it as a strange or misguided way to investigate and understand something as practical as morality. It will thus be worth explaining and justifying my chosen starting points, and I will do this in §1.2 and §1.3. First, however, I will clarify what Robust Realism actually is.

1.1 Robust Realism

It's often said that 'realism' is a term of art, and this is certainly true in metaethics.² Some philosophers have attempted to provide a principled distinction between realism and anti-realism in the metaethical domain, but given the technical nature of these terms of art I am mostly content to be permissive about who gets to call themselves a realist about morality.³ The 'realism' label is apt for a variety of metaethical positions, and at the moment I feel no particular push to offer even a stipulative definition. As long as we are able to distinguish between different metaethical theories by identifying the specific cluster of commitments that each such theory accepts, I'll be happy enough. My goal in this section is thus to (briefly) outline the exact combination of commitments that characterises Robust Realism, and that marks it out as distinct from various rival metaethical views.

An initial commitment that all Robust Realists will accept is as follows:

Truth. There are substantive moral truths.

Candidates for such truths include claims like 'it is wrong to kill for fun' and 'helping others is generally admirable.' Note that the commitment is to *substantive* moral truths. Everyone can endorse a claim like 'killing is wrong or nineteen is a prime number,' or 'if killing is wrong, then killing is wrong.' Even if such propositions are moral truths, they are too trivial to be of interest. However, whilst Truth is a claim that all Robust Realists will accept, it is not nearly enough to carve out distinctive territory for them.

² Cf. Wright (1992: 1).

³ For attempts to provide a principled distinction, see Dreier (2004) and Miller (2009).

This is because Truth can be interpreted in different ways, and many of those who style themselves as moral anti-realists will also claim to capture Truth. For instance, self-styled moral anti-realists might seek to accommodate substantive moral truth by appealing to deflationary, or relativist, or constructivist accounts of such truths.⁴ As we'll see, there are a number of additional commitments that distinguish a *robustly* realistic interpretation of Truth from both realist and anti-realist alternatives. The overall combination of commitments that Robust Realists defend marks them out as offering an unusually uncompromising conception of moral reality.

In particular, Robust Realists think that we need a metaphysically committed understanding of the claim that there are substantive moral truths. They will therefore accept the following:

Properties. At least some substantive moral truths pertain to the moral properties of certain acts, or types of act.

This puts us on the right track, but to get to a metaphysically robust form of moral realism we'll need to say more about the nature of these moral properties. As it stands, the commitment to Properties is compatible with several theories that are not as uncompromising as Robust Realism. It is compatible with theories according to which moral properties reduce to descriptive properties, theories according to which moral properties are within the purview of the sciences, and theories according to which moral properties are in some way subjective. Robust Realists see each of these views as unsatisfactory, for reasons discussed in later chapters. So, to get to Robust Realism, we must clarify just what is involved in rejecting these views.

Robust Realists reject reductive construals of moral properties, and instead accept the following:

Non-Reductivism. Moral properties are not reducible to purely descriptive properties.

But what are these 'purely descriptive' properties, and what is it for moral properties to be (or fail to be) reducible to such properties? I return to these questions in Chapter 3, but it is worth saying just a little about them now.

⁴ For a deflationary understanding of substantive moral truths, see Blackburn (1998). For a relativist understanding, see Harman (1977). For a constructivist understanding, see Street (2008; 2010).

Descriptive properties are those properties that can be picked out by the descriptive predicates, which include predicates like ‘... is human’ and ‘... is a garden gnome.’ Now, this gloss raises an obvious question: what is a descriptive predicate? And this question has an obvious answer: a descriptive predicate is one that describes something as it is or can be. Typically, then, to use a descriptive predicate is to say something true or false. This might make Non-Reductivism seem odd. After all, there is an everyday sense in which to say that an act is right or good is to describe it as right or good. And moral claims can be true or false.⁵ So is it not clear that moral properties are descriptive? It is not. At least, it’s not clear that they are *purely* descriptive, for moral predicates like ‘... is right’ involve or entail normative direction. This means that they’re not just capable of truth or falsity, for directions can also being satisfied or unsatisfied. This directive element is hard to see in descriptive predicates like ‘... is human’ and ‘... is a garden gnome,’ or even ‘... is desired’ and ‘... is pleasurable.’ These appear to be entirely on the ‘is’ side of the ‘is/ought’ divide. They seem to lack the element of direction found in morality. This appearance might turn out to be illusory, but Non-Reductivism amounts to a denial of the claim that the directive bit of morality can be had in purely descriptive terms.⁶

What about reduction itself? In this context, we can see reduction as a relation that holds between properties or classes of property. There are various ways of understanding what it is for one class of properties to reduce to another. For instance, one might view the reduction relation as the identity relation, or as the constitution relation, or as the grounding relation. Alternatively, one might avoid a specific approach altogether by using a generic placeholder. For instance, one might ask whether moral properties exist ‘over-and-above’ purely descriptive properties, or one might ask whether they can be said to constitute an ‘addition of being.’⁷ The approach to reduction that will be relevant in Chapter 3 involves the identity relation. I understand Robust Realism as primarily concerned to deny that the moral properties can be identified with purely descriptive properties.

⁵ Moreover, *thick* moral predicates – ‘... is brave’ and ‘... is cruel’ – have descriptive content.

⁶ The opposing view – ‘Descriptivism’ – is defended by Jackson (1998) and Brown (2011). I discuss this position in depth in Chapter 3.

⁷ Enoch (2011: 101) speaks of an ‘over-and-above’ relation. The phrase ‘addition of being’ is from Armstrong (1997: 117), though I don’t know if he is using it as a placeholder.

Robust Realists don't just deny that moral properties reduce to any purely descriptive properties. They also deny that moral reality can be naturalised.⁸ It is hard to say exactly what this means, for the character of the natural/non-natural distinction is controversial in metaethics, and indeed philosophy generally. However, as I explain in Chapter 5, I operate with a 'methodological' approach to the natural/non-natural distinction. More specifically, I will work with a 'disciplinary' view of this distinction, on which the natural properties are those that fall within the purview of the natural or social sciences. So, a property counts as natural if it can be discovered by the methods of disciplines like physics, chemistry, biology, neuroscience, psychology, economics, sociology, anthropology, and so on.⁹

This sort of methodological approach is common in metaethics, but it does raise questions. For instance, if we view natural properties via the natural and social sciences, we might want a more detailed account of what it is for some discipline to count as a natural or social science. And this complicates matters, for even if most of us know a science when we see one, demarcating science from non-science won't be straightforward. Still, even if we're unsure of precisely which conditions are necessary and sufficient for something's counting as a science, we can plausibly suppose that there are such conditions. For instance, it is likely that science, unlike non-science, tests its hypotheses using repeatable empirical experiments. There is more to science than this, as we'll see in Chapter 5, but there are things that we can say to usefully distinguish scientific and non-scientific inquiry.

One might also wonder if it is too artificial to divide scientific from non-scientific inquiry. Universities often put their science and philosophy departments in separate faculties, but this hardly entails that they lack methodological similarities. I've no doubt that there is a certain overlap (see below). But there are also key differences, as I explain in Chapter 5. Thus, if we treat the disciplinary approach with care, and don't take too much from it, it is a neat way to approach the natural and the non-natural. I know of no better way of approaching it, so I'll put this one to cautious use.

⁸ For reasons that will become clear in later chapters, we should distinguish the rejection of reductive theories from the rejection of naturalistic theories. This is basically because not all naturalistic theories are reductive, and not all reductive theories are naturalistic.

⁹ For detailed discussion, see Copp (2003), Shafer-Landau (2006), and Cuneo (2007b).

In short, Robust Realists reject naturalistic metaethical theories and instead accept the following:

Non-Naturalism. Moral properties are beyond the purview of the natural and social sciences.

I will consider how Robust Realists can best interpret and motivate the commitment to Non-Naturalism in Chapter 5, but for now the point is just that this commitment is sufficient to distinguish their theory from those that seek to accommodate morality within a strictly scientific picture of reality.¹⁰

In addition to defending the irreducibility and the non-natural status of moral properties, Robust Realists also say that these properties exist objectively. There are different ways in which something might be considered objective. For instance, one might say that an entity counts as objective if one can be mistaken about its existence or nature. However, the Robust Realist has a much stronger style of objectivity in mind. This is because one can be mistaken about the existence and nature of socially constructed entities like money and nation states, and the Robust Realist certainly does not think that moral properties or truths are entities of this sort.¹¹

Robust Realists instead accept the following very strong view of moral objectivity:

Mind-Independence. Moral properties are constitutively independent of any agent's or set of agents' actual or hypothetical attitudes, beliefs, and conventions.

This does not entail that our attitudes (and so on) are entirely irrelevant to what is morally right or wrong in particular cases. For instance, it is compatible with Mind-Independence that substantive moral obligations lie solely in doing as one desires. I'd be surprised if this is where all such value lies, but this claim is perfectly compatible with Mind-Independence. If one's obligation does indeed lie solely in doing as one desires, defenders of Mind-Independence will just say that this is the case even if you judge that one should never do as one desires and that just having desires is evil. This is the strong form of moral objectivity to which Robust Realists commit, and I discuss this commitment in detail in Chapter 4 and Chapter 6.

¹⁰ For some naturalistic theories, see Boyd (1988), Brink (1989) and Thomson (2008).

¹¹ This distinguishes them from constructivists like Korsgaard (1996) and Street (2008).

So far I have said that Robust Realists regard moral properties as irreducible, non-natural, and mind-independent. Now, these are claims about the *existential* status of such properties, but the Robust Realist also makes a claim about the *authority* of morality:

Categorical Authority. The authority of moral direction is paradigmatically categorical.

First off, what is it for moral direction to be categorical? I will answer this question carefully in Chapter 5, but for now we can just say that direction is categorical if it is unconditional, or inescapably binding.¹² Categorical direction is independent of any contingencies of a social or psychological nature.¹³

We can contrast categorical direction with hypothetical direction. Consider the following: *if you want you impress your boss, you should offer to work late this evening.* There is certainly direction here, for the claim that you should offer to work late requires or demands something of you. But there is no categorical direction in this claim, for it is conditional on a particular psychological contingency – wanting to impress your boss. And the fact that this piece of direction comes with conditions ensures that it is escapable. If you were to change your end (impressing your boss), then (assuming no other relevant conditions are in play) the direction to take the means to that end (offering to work late) doesn't apply. Hypothetical direction is thus binding, but only for those with the relevant desires or ends.

This means that hypothetical direction is distinct from categorical direction like the following: *you ought to keep your promise to meet your friend.* It might get trumped by another moral obligation, but you cannot escape this piece of moral direction just because you don't want to meet your friend. The obligation to keep your promise is unconditional, and is thus in some sense inescapable. The same goes for most other pieces of moral direction, at least according to Robust Realists. I will expand on this commitment in Chapter 5, offering a more detailed approach to morality's inescapability and the role that this can play in attempts to motivate the rejection of naturalistic theories. Before moving on, it is worth drawing out two points about the Robust Realist's commitment to Categorical Authority.

¹² Cf. Kant (*Groundwork*).

¹³ This way of putting it echoes Hampton (1998: 96).

First, the commitment does not require that *all* moral direction is categorical. Just that it is paradigmatically so. This means that, if we were to discover that some moral direction is non-categorical, this would not be sufficient to falsify Categorical Authority. And this is perfectly fine, for Robust Realism should be understood as a *particular* thesis rather than as a *universal* thesis. In other words, it is a theory on which *some* moral properties are robustly real, not on which *all* such properties are robustly real.

Second, note that the commitments to Categorical Authority and Mind-Independence are distinct. I want to make this explicit, for they look similar. After all, both are independence claims. Mind-Independence says that moral properties are constitutively independent of attitudes, beliefs, and conventions. Categorical Authority says that the authority of moral direction is independent of social and psychological contingencies. There is a difference between these independence claims. Mind-Independence is a claim about the objective *existential* status of moral properties, whilst Categorical Authority is a claim about the inescapable *authority* of the direction provided by the ascription of such properties. If the difference remains unclear, consider that it is entirely uncontroversial to hold that natural properties can be mind-independent, but it is highly controversial to hold that the ascription of natural properties can direct categorically. Maybe categoricity can be naturalised, but this is something that must be shown.

Before defining Robust Realism, there is another point to address. Some have argued that the commitments discussed so far can be secured without any robust metaphysical commitments. This claim is associated with two theories – ‘Quasi-Realism’ and ‘Relaxed Realism.’¹⁴ I’ll discuss these theories in the next chapter, but the idea that they have in common is as follows: when the commitments that we have considered so far are correctly construed, they are *moral* and not metaphysical commitments. They are ‘internal’ to ethics, occurring within the moral domain. They are not ‘external’ claims *about* ethics, requiring a metaphysically committed theory. Indeed, the Quasi-Realists and the Relaxed Realists have typically suggested that these commitments can *only* be understood as internal to ethics.

¹⁴ Quasi-Realists include Blackburn (1984; 1993a; 1998; 2010a) and Gibbard (2003). Relaxed Realists, also known as ‘Quietists,’ include Dworkin (1996; 2011), Nagel (1997), Kramer (2009), Parfit (2011 v.2), and Scanlon (2014).

If this is correct, then there is no space for a metaphysically robust view of moral reality. Of course, Robust Realists think that there is space for such a view. They therefore accept the following:

External Metaphysics. We can assess the nature of moral properties from an external metaphysical standpoint, and the existence of such properties is to be defended or rejected primarily from that standpoint by appeal to external metaphysical considerations.

I will explain and defend this commitment more thoroughly in the next chapter, showing how Quasi-Realism and Relaxed Realism fail to threaten it. For now we can just note that a commitment to External Metaphysics is required to demarcate Robust Realism from any theory that claims to secure claims like Truth, Properties, Non-Reductivism, Non-Naturalism, Mind-Independence and Categorical Authority in some non-metaphysical sense.

Now we can define Robust Realism. The following is sufficient to characterise the heart of the theory, and to distinguish it from rivals:

Robust Realism. There are substantive moral truths that commit us to irreducible, non-natural, mind-independent moral properties. The authority of direction provided by the successful ascription of such properties is paradigmatically categorical. Their existence is to be defended primarily from an external metaphysical standpoint, by appeal to external metaphysical considerations.

This has been a whistle-stop tour through the terrain that Robust Realism occupies in the metaethical landscape. I won't go into further detail about the commitments that characterise my conception of Robust Realism, for I return to them all in subsequent chapters. What matters now is that this constitutes an operable definition of the metaethical theory that I wish to examine.

1.2 Metaphysics

You might have noticed that Robust Realism, as I've defined it, is a strictly metaphysical theory. It incorporates a number of claims about the nature of moral reality, but it incorporates no claims about moral semantics or psychology. I will explain why I define Robust Realism in this way in §1.3, but first I will say more about metaphysics and why I take it as my starting point.

After all, you might see it as a mistake to do so. Perhaps you are suspicious of metaphysics in general, or perhaps you think that abstract metaphysical theorising is a peculiar way to start an investigation into as blatantly practical a domain as morality. I agree that morality is primarily a practical discipline. For the most part, moral life does not take place in a seminar room. Without wishing to offer any sort of definition, moral inquiry is centrally concerned with how we should act and how we should live. In contrast, metaphysics seems extremely abstract and theoretical. Without wishing to offer any sort of definition, metaphysics is centrally concerned with understanding the nature and structure of reality as a whole.

More specifically, as I understand it, metaphysics is an inherently speculative enterprise. A metaphysician considers the various ways that reality *might* be, and then gives tentative arguments for thinking that it *is* one way and not another. These needn't be knockdown arguments. That would be an unfairly high expectation, and metaphysical inquiry can be fruitfully conducted with much more humble intellectual ambitions. And this means that metaphysicians need not retreat to studying the structure of our concepts, or whatever. Given the caveat that their suggestions are speculative, they can view themselves as studying the nature of reality itself.¹⁵

Still, you may ask why one should *start* a metaethical inquiry with metaphysics. Given that morality is primarily practical, and given that metaphysics seems primarily theoretical, what bearing can metaphysical inquiry have on our view of morality? The first thing to note here is that it is simply a mistake to operate with a strict division between theory and practice. The distinction may sometimes be useful, but it is also artificial. Theory can matter for practice, and practice is at least to some degree theory-laden. Our engagement with practical issues takes place against a background of internalised concepts, norms, motives, rationales, and so on. These are often related more or less systematically and more or less coherently, and are thus more or less theoretical. They may not rise to the level of explicit and unified philosophical theories, but they can still be understood as theories that guide practice. One job for philosophers is to make these theories explicit, to assess their quality, and to improve upon them.¹⁶

¹⁵ For this conception of metaphysics, see Lowe (1998; 2000).

¹⁶ For a more detailed discussion in a similar vein, see Nussbaum (2000).

This point applies even with highly practical domains. Answers to substantive ethical questions about the permissibility of euthanasia, for example, will turn partly on metaphysical assumptions about the nature of personal identity. And those metaphysical assumptions may turn upon on further and perhaps much more abstract metaphysical assumptions about, for example, agency and the identity relation. This is certainly not to say that the metaphysical issues must be settled before one can discuss the ethical issues. It's simply to say that metaphysical assumptions often bear on our answers to moral questions. One role for theory is to shape these background assumptions into something coherent and plausible, so metaphysical theorising has a contribution to make to substantive moral inquiry.

The insights provided by moral metaphysics will also have some practical significance. Metaphysical argument may ultimately convince us that there are no moral properties, for example, and this is very likely to have an effect on our ability to take seriously those claims that involve such properties. We may feel compelled to abandon morality, or to revisit and perhaps rethink our understanding of what morality is. Alternatively, metaphysical inquiry might convince that there are moral properties, and it might help us to understand them. The discovery that there are moral properties would vindicate morality against certain sceptical attacks, and could help to reinforce our efforts to work together on moral questions. So, once again, metaphysical theorising is not without practical import. It is thus a mistake to reject metaphysical inquiry as a metaethical starting point on the grounds that, unlike moral life, it is primarily theoretical and abstract.

Moreover, the basic concern of metaphysics ensures that it will come up eventually in any prolonged intellectual inquiry. Metaphysics is concerned with scrutinising the nature and structure of reality as a whole. Its province is thus all-encompassing. It is worth noting that this inclusive conception of metaphysical inquiry entails that even scientists do a little bit of metaphysical (or at least quasi-metaphysical) inquiry at a particular point in their investigations. When they offer their hypotheses, scientists are speculating as to how a certain portion of reality might be. Of course, scientists seek empirical evidence (rather than philosophical arguments) for their hypotheses – this is one crucial way in which scientific inquiry diverges from that of philosophical metaphysics, as I'll explain in Chapter 5. Nevertheless, there is a certain level of overlap between these forms of inquiry.

On this approach, a little bit of metaphysical inquiry is therefore unavoidable even if science is the only path to knowledge. But one should not assume without argument that science is the only path to knowledge. Indeed, as I noted in §1.1, the Robust Realist accepts Non-Naturalism, the view that the moral properties are beyond the purview of the natural and social sciences. Either way, to inquire into the way in which morality fits into the universe is to consider various ways that moral reality might be, and to then defend one particular conception of it. For the Robust Realist, then, inquiry into the nature of moral reality is an inherently speculative enterprise. And that's fine. Not all speculation is idle, even if it is backed up by philosophical arguments rather than empirical experiments. But, again, what reasons are there to *start* with this speculative metaphysical inquiry.

Well, as should now be obvious, it's going to come up eventually given the all-encompassing nature of metaphysical inquiry. Metaphysics is centrally concerned with the nature and structure of reality as a whole, and moral life is a part of the nature and structure of reality as a whole. Anyone interested in giving a systematic understanding of morality will thus have to do a bit of metaphysical inquiry at some point. Bluntly, then, one might as well get it over with. Less bluntly, in asking about the nature of ethics and ethical life, it seems to me sensible to start with the broadest form of inquiry. Understanding and comparing the various ways that moral reality might be will give frameworks within which other questions – for instance, questions about moral semantics and psychology – can be located.

I lack the space (and the energy) to develop and compare every account of moral reality, so I centre my discussion on one theory – Robust Realism. I don't aim to show that this theory is true, but I will develop a compelling metaphysics for it. Having now offered some general remarks about metaphysics and its legitimacy as a metaethical starting point, I will next explain the particular significance of metaphysical inquiry to Robust Realism.

1.3 Metaphysics and Robust Realism

I have been explicit about the fact that I define Robust Realism in strictly metaphysical terms. This is not because I regard the various semantic and psychological claims that are commonly associated with Robust Realism as unimportant. Rather, it is because I do not see them as *essential* to this theory.

First, what semantic and psychological claims are associated with Robust Realism? We can start with the following:

Cognitivist Semantics. Moral sentences are truth-apt.

Cognitivist Psychology. Moral judgements express beliefs.

Of course, a great many non-robust metaethicists will also accept these rather thin commitments, but Robust Realists can distinguish themselves by accepting the following:

Robust Referent. If moral sentences and moral judgements successfully refer, they refer to robust moral properties.

This is an interesting and controversial set of commitments, but it is not essential to Robust Realism. We can see this by considering the possibility of combining a robust account of moral properties with a substantive view known as ‘Abolitionism.’ As will become clear, this is an unorthodox pairing of views, but its mere availability within the conceptual landscape shows that Robust Realism is best seen as a fundamentally metaphysical theory.

In this context, Abolitionism is the view that it is in our interests to eliminate moral thought and discourse. Its defenders recommend that we stop uttering moral sentences and making moral judgements. They allow that we can engage in non-moral direction and evaluation, but think that we are better off getting rid of our moral framework. Abolitionism is associated with ‘Error Theory,’ but it is also possible to combine it with other metaethical views. Before we consider how it can be combined with Robust Realism, it will help to see how it relates to Error Theory. First off, note that Error Theorists are with Robust Realists when it comes to moral semantics and psychology. They part, however, in their willingness to accept the metaphysical views outlined in §1.1 – whereas Robust Realists say that robust moral properties exist, Error Theorists deny this. In short, Error Theorists argue that moral thought and discourse are essentially committed to the existence of robust moral properties, that no such moral properties exist, and that moral thought and discourse are thus infected with an error.¹⁷ If this is correct, an obvious question arises: what are we to do? Abolitionism offers one possible answer. It tells us that we are to get rid of our moral framework, and direct or evaluate only in non-moral terms.

¹⁷ Cf. Mackie (1977), Joyce (2001; 2006), Streumer (2008; 2011), and Olson (2014).

This is not, of course, the only available option. According to some Error Theorists, it is in our interests to reform our moral framework so that it becomes a sort of make-believe. This is 'Fictionalism.' On this view, rather than making moral judgements we would instead just *pretend* to make moral judgement, for there is no error in pretending. Alternatively, we might accept 'Conservationism.' On this view, we should just continue with our faulty moral discourse.¹⁸ However, the Abolitionist recommends that we eliminate our moral framework altogether.¹⁹ Those who defend this view do so because they think that, prudentially speaking, it is our best option. They maintain that moral thought and discourse do badly in a cost-benefit analysis, and that it is in our interests to replace them with something else – probably something like purely prudential thought and discourse. This will make life go better, whilst avoiding any metaphysical issues.

Various arguments for Abolitionism have been offered, but the most effective centres on the categoricity of moral direction. Roughly, the suggestion is that categorical moral judgements undermine our ability to successfully resolve practical conflicts. More specifically, if you judge that Φ -ing is unconditionally and inescapably demanded of you, this will make you less inclined to compromise with those who are opposed to Φ -ing. The risk is that this will lead to entrenchment and deadlock among the conflicting parties in difficult disputes, making it less likely that they will productively resolve their practical conflicts at both an interpersonal and an intercultural level. Whether or not you find this plausible, it is clear that Abolitionism is a view that can be coherently combined with Error Theory. What is of interest here, however, is that Abolitionism can also be coherently combined with Robust Realism.²⁰ If it is in our interests to abolish our moral framework, a Robust Realist can recommend that we do so. Indeed, unlike the Error Theorist, the Robust Realist can even claim that our moral framework should be abolished for moral reasons; maybe the most effective way of realising moral ends is to stop thinking in moral terms.

¹⁸ Fictionalism has been defended by Joyce (2001; 2005) and Nolan et al. (2005), whilst Conservationism has been defended by Olson (2014).

¹⁹ Abolitionism is defended by Hinckfuss (1987), Garner (2007), and Marks (2013). I have argued that this theory has a surprising amount going for it in Ingram (2015a).

²⁰ I think that Robust Realism and Abolitionism make an interesting theoretical package. Although I do not in fact endorse this package, I have explored it in Ingram (2015a).

If we were to stop thinking in moral terms, this wouldn't mean that robust moral properties don't exist. They're mind-independent, so they'll still exist even if you aren't paying attention to them. We might not even have to eliminate the second-order belief that such properties exist. We just have to stop appealing to them in ordinary life, and replace our moral framework with something else. In practice, this would of course be difficult. But it is at least possible in principle, and the mere possibility of combining a robust view of moral properties with Abolitionism shows that we should understand Robust Realism as an essentially metaphysical view.

Let's name the combination of Robust Realism and Abolitionism 'Realist Abolitionism.' It would be odd to say that Realist Abolitionists fail to be robust in their metaethical outlook even though they maintain that there exist irreducible, non-natural, mind-independent, and categorically authoritative moral properties whose existence can be established from an external metaphysical standpoint. The fact that Realist Abolitionists also take there to be compelling moral or prudential reasons to abolish moral thought and discourse should not prevent us from seeing them as metaethically robust. The central insight of the Realist Abolitionist is that, if there are robust moral truths, they're not going to go away just because the final moral judgement has been made, and the final moral sentence uttered.

This is why metaphysical inquiry should be our starting point in an investigation into Robust Realism.²¹ Semantics and psychology are both interesting, and they may help when we think about the more basic metaphysical issues. Still, it is the metaphysical commitments outlined in §1.1 that constitute the essential core of Robust Realism. It is with these commitments that Robust Realism must stand or fall, and metaphysical considerations will thus be vital to our evaluation of this theory. Even if we set aside the general reasons to begin metaethics with metaphysical inquiry, there are reasons for inquiry into Robust Realism to start in this way.

²¹ Kahane (2013) arrives at a similar conclusion in a different way. He argues that we can see that Robust Realism is fundamentally metaphysical by inverting Error Theory. After all, Error Theory accepts a robust moral semantics and psychology, but rejects a robust metaphysics. Kahane discusses so-called 'Reverse Theory,' on which a robust metaphysics is accepted but a robust semantics and psychology are rejected. He plausibly argues that the possibility of Reverse Theory shows Robust Realism to be essentially metaphysical.

1.4 Necessities

Having explained the importance of metaphysics for investigating the nature of morality in general and Robust Realism in particular, I am now in a position to lay bare something that will be important in subsequent chapters but that has hitherto gone unmentioned. Specifically, a key idea that guides the discussion in subsequent chapters is that considerations pertaining to the *necessity* of (at least some) moral truths are crucial for understanding and motivating the metaphysical commitments of Robust Realism.

For example, some influential objections to Non-Reductivism can be answered by appeal to the idea that there are necessary moral truths. Moreover, a plausible account of these necessary moral truths is, I will argue, partly bound up with the commitment to Mind-Independence. In addition, the commitment to Categorical Authority can be interpreted via a certain sort of necessity claim, and this in turn helps the Robust Realist to motivate their commitment to Non-Naturalism. All of this will become clearer as the discussion unfolds, and I will not pre-empt that discussion at this point. However, given that considerations pertaining to necessity will have an important role to play in subsequent chapters, it is worth taking a moment to provide some general comments about necessity and modality.

First of all, what is necessity? The short answer, which will do for my purposes, is that something is necessary if it *must* be. In other words, the necessary is that which cannot be otherwise. For instance, it is clearly necessary that $2+2=4$. This cannot be otherwise. In contrast, something is contingent (or merely possible) if it only *might* be. This is to say that the contingent is that which *can* be otherwise. It is merely contingent, for example, that I have curly hair. I could have had straight hair, or I could have had no hair at all. Now, claims about necessity and contingency are often articulated using the language of 'possible worlds,' and I follow this practice. Roughly speaking, possible worlds are complete ways the actual world might have been.²² Setting aside a caveat mentioned below, we can say that a proposition is a necessary truth if it is true at all worlds, and we can say that it is a contingent truth if it is true at some worlds and not at others.

²² As far as I can tell, nothing that I say will depend on any particular conception of the nature and status of possible worlds. For some prominent views, see Plantinga (1974), Kripke (1980), Lewis (1986), and Armstrong (1989).

To clarify, in saying that it is necessary that $2+2=4$, one is saying that this is true at all possible worlds. In saying that it is contingent that I have curly hair, I am saying that there are worlds in which I do not have curly hair. In some of these worlds, I do not exist. In others, I exist with straight hair. In others, I exist but have no hair at all.²³ And so on. As will become clear in later chapters, the Robust Realist can achieve a great deal just by elaborating on the claim that some moral truths hold at all possible worlds.

Now, it may be that there is more than one type of necessity. We can speak of analytic necessity, nomic necessity, metaphysical necessity, epistemic necessity, conventional necessity, and so on. Whether there is a deep distinction between these forms of necessity is up for debate, but we can differentiate them in a somewhat loose manner.²⁴ For instance, a proposition is analytically necessary if it must be true given its meaning.²⁵ It is nomically necessary if it must be true given the laws that govern the relevant phenomena. And it is metaphysically necessary if it must be true, even under different laws of nature. Controversies abound when it comes to giving a detailed account of these necessities, and the relations among them all, but my only concern in what follows is the necessity in Robust Realism.

Note as well that necessity claims can be more or less restricted in scope. Not all uses of modal terms like 'must' and 'cannot' pertain to all worlds. Some are implicitly or explicitly localised to a world or a set of worlds. Consider tennis. To win a tie-break game a player *must* reach at least seven points, with a margin of two points over the opposing player. However, the scope of this 'must' is (implicitly) restricted to the worlds at which certain conventions – the rules of tennis – obtain. And there are worlds in these conventions are different. For instance, there are worlds in which a player must reach at least nine points to win a tie-break. The nature of the conventions limits the range of worlds in which the claim is true.²⁶

²³ I'd rather not think about such worlds, but I am willing to acknowledge their possibility.

²⁴ For discussion of how the varieties of necessity may be related, see Fine (2002).

²⁵ Of course, there are complications arising from Quine (1951).

²⁶ Most games have some rules analytically built into them. To change such rules would be to make a new and different game. I don't think that the tie-break rule in tennis is such a rule. But it is likely true of other rules in tennis, like the rule requiring a player to hit the ball so that it lands within the court. Thanks to Jimmy Lenman for this observation.

Although many necessity claims are restricted in their scope, for the Robust Realist it is mainly those that extend to all possible worlds that will be at issue. Again, this is because the Robust Realist can go a long way simply with the suggestion that certain moral truths obtain at all possible worlds. But how should we understand the necessity in Robust Realism itself? This is an important and interesting question that is only beginning to receive the attention it deserves.²⁷ I discuss it in Chapter 4, in which I tentatively develop a comprehensive account of the necessity in Robust Realism.

1.5 Spoilers

In this chapter I have explained and justified my starting points. I begin with moral metaphysics, which I take to be the speculative enterprise of outlining and assessing the various ways that moral reality might be. My specific focus is on Robust Realism, the claim that there are irreducible, non-natural, mind-independent, and categorically authoritative moral properties that exist from an external metaphysical point of view. In the chapters that follow, my overarching aim is to develop this theory in more depth. I mean to outline a unified and compelling metaphysics for Robust Realism. Moreover, in developing this metaphysics I will at various points draw on claims relating to necessity. However, before launching into this project, I will first explain how exactly it is going to unfold. I don't like to give the game away, but it may help to have a sense of what each chapter will do.

Chapter 2 addresses some challenges to External Metaphysics. It is important to begin with this commitment, as there will be no space for a robust view of moral properties if it turns out to be untenable, and that would obviously undermine my core aim, which is to outline a compelling metaphysical position for Robust Realism. I explain the precise nature of the challenge facing defenders of External Metaphysics, distinguishing a modest and an ambitious version of this challenge. The modest version is associated with Quasi-Realism, and the ambitious version is associated with Relaxed Realism. My discussion of these theories will show that we can uphold a commitment to External Metaphysics, properly interpreted, and thereby vindicates my project of providing a metaphysics for Robust Realism.

²⁷ A recent discussion comes from Cuneo and Shafer-Landau (2014). I consider and reject their view in Chapter 4. For other discussions, see Scanlon (2014) and Skarsaune (2015).

Chapter 3 examines the commitment to Non-Reductivism. After outlining some evidence for Non-Reductivism, I consider two influential objections to it. I show that these objections can be answered by appeal to necessary moral truths. The objections in question seek to challenge Non-Reductivism by appealing to the supervenience of moral properties on descriptive properties. The first says that this supervenience relation, combined with some other plausible claims, entails that moral properties reduce to descriptive properties. The second says that Non-Reductivism makes a mystery out moral supervenience. I explain how both objections can be answered by appeal to necessary moral truths. I do not argue that there definitely are such truths, but I do argue that this claim is at least tenable.

Chapter 4 examines the necessity in Robust Realism in greater detail. I consider and reject a recently suggested view according to which some moral truths hold as a matter of conceptual necessity. I then outline a different account of the necessity in Robust Realism. I basically suggest that the necessity in Robust Realism can be understood in terms of the elimination of what I call 'limiting factors.' These are, broadly, contingent factors that are built into or presupposed by many moral norms, and that limit the range of worlds in which those norms have authority. If a moral norm can survive its limiting factors having been stripped away, and if all of these limiting factors are in fact stripped away, we will have found a moral norm that holds as a matter of necessity. I illustrate this process of stripping away limiting factors with Mind-Independence, a metaphysical commitment which in itself eliminates a significant limitation on moral norms.

Chapter 5 looks at the commitment to Categorical Authority, and explores how it can help to motivate the commitment to Non-Naturalism. Interpretations of categoricity often rely on evocative but vague images, so I provide a way of getting a clear grip on it. Roughly, we can get a grip on the notion of categorical authority by focusing on the fact that anyone who culpably violates a categorical moral norm will necessarily acquire the property of being morally blameworthy. This necessity claim can then be construed in a modally 'fragile' or a modally 'robust' sense. If there are necessary moral norms then the modally robust sense is the way to go, and I argue that it leads us to Non-Naturalism. More precisely, I argue that the modally robust necessitation of moral blameworthiness is something that even the most promising version of moral naturalism is unlikely to capture.

Having developed this set of metaphysical positions, and having explained why Robust Realists should find it compelling, I then move on to a different topic. Chapter 6 examines the dialectical situation that we face in contemporary metaethics. I consider this in order to assess how much we can expect to achieve when we argue about moral reality. Given that my aim here is just to develop Robust Realism, rather than to show that it is true, I want to know whether and how we can make progress in metaethical debate. Building my discussion around Mind-Independence, I suggest that some common ways of arguing about moral reality lead to an intractable stalemate, and that this stalemate is likely to generalise and persist as a result of entrenched differences in our temperaments and existential needs. I explain how these temperamental divisions have a lot to do with our feelings about the link between necessity, contingency and arbitrariness. I then consider how one might respond to this dialectical situation, arguing that we may need to limit the ambitions of metaethical inquiry. In spite of our best efforts, determining whether Robust Realism (or some other metaethical theory) is actually *true* might well be beyond us.

Chapter 7 concludes my inquiry by bringing together its various elements, and by sketching a number of avenues that I hope to explore in future research. However, before thinking about research that I hope to do in the future, I'd better do some in the present. That starts in the next chapter.

2 The Possibility of Robust Realism

Robust Realism is distinguished from most other metaethical theories by its acceptance of the following combination of commitments:

Truth. There are substantive moral truths.

Properties. At least some substantive moral truths pertain to the moral properties of certain acts, or types of act.

Non-Reductivism. Moral properties are not reducible to purely descriptive properties.

Non-Naturalism. Moral properties are beyond the purview of the natural and social sciences.

Mind-Independence. Moral properties are constitutively independent of any agent's or set of agents' actual or hypothetical attitudes, beliefs, and conventions.

Categorical Authority. The authority of moral direction is paradigmatically categorical.

It will be useful to give this combination of commitments a name, so I will call them the 'Realist Commitments.' Given that 'realism' is a term of art, you can legitimately describe yourself as a moral realist even if you don't accept each and every one of the Realist Commitments. However, they are all in fact accepted by each of the three theories that are at issue in this chapter.

In addition to Robust Realism, the theories at issue in this chapter are Quasi-Realism and Relaxed Realism. There are important differences between those two theories, but there are important similarities too, for they are roughly alike in how they see the Realist Commitments. Both see them as *moral* rather than *metaphysical* commitments, and both maintain that they can be secured without any robust metaphysical claims. On top of this, Quasi-Realists and Relaxed Realists commonly suggest that the Realist Commitments can *only* be interpreted in this non-robust light. Put another way, for many Quasi-Realists and Relaxed Realists, there is just no space for a metaphysically robust reading of the Realist Commitments. This should clearly worry anyone with plans to develop a metaphysically robust interpretation of the Realist Commitments. Given that I have plans to do just that, I first want to answer the Quasi-Realist and the Relaxed Realist.

After all, if the Realist Commitments cannot be interpreted as metaphysical, it won't even be possible to articulate (let alone defend) a robust view of morality. This is the sense in which many Quasi-Realists and Relaxed Realists have denied that there is space for Robust Realism. So, these two theories present an apparent threat to the very *possibility* of Robust Realism. In this chapter I deal with this threat, explaining what it is and how it fails. I begin in §2.1 and §2.2 by outlining the challenge in detail, exposing a modest and an ambitious version of it. Quasi-Realism presents the modest version, and Relaxed Realism presents the ambitious version. I discuss Quasi-Realism in §2.3 and §2.4, showing that it fails to threaten the possibility of Robust Realism. I then move on to the chapter's main focus, which is the more ambitious challenge defended by Relaxed Realists. I explain the Relaxed Realist position in §2.5 and §2.6, and I then respond to the threat that it poses in §2.7 and §2.8. I conclude in §2.9 that there is space to articulate and defend a metaphysically robust account of morality.

2.1 Two Points of View

What goes into a robustly metaphysical view of the Realist Commitments, and how could one view them in any other way? We can begin to answer these questions by appealing to a distinction between the two points of view from which one can think about morality. On the one hand, we have the 'internal' point of view. On the other, we have the 'external' point of view.

First, the internal point of view. This is the first-order standpoint from which we ask and answer substantive ethical questions. It is called 'internal' because it is the point of view occupied *within* ethics. In other words, it is the perspective from which one judges that killing is wrong, that helping other people is admirable, and that courage is a virtue. These claims are all quite widely accepted, of course, but we can also ask more controversial questions at the internal moral level. For instance, we might assess whether it is morally permissible to eat meat by considering moral arguments for and against meat eating. More abstractly, we might ask what it takes for an act to be right, for a person to be good, and for a social institution to count as just. We can work on these first-order moral issues together, for the internal perspective is one that all moral agents occupy. Of course, we often disagree about the best way to answer substantive moral questions. These are moral disagreements, for they occur within ethics.

Second, we have the external point of view. This is a second-order standpoint at which we abstract away from all straightforwardly ethical questions and ask *metaethical* questions instead – that is, questions about moral reality, moral semantics, moral psychology, moral knowledge, and so on. This external perspective is often seen as being ‘detached’ from the internal moral perspective.¹ On this understanding, when we occupy the external point of view can be strictly neutral on the substantive moral questions that we ask within the internal perspective. The external point of view is therefore distinctively non-substantive. Put another way, every external metaethical position is compatible with every first-order moral position.

I’ll say more about this later, but for now the point is that it is from an apparently neutral perspective that one asks metaethical questions. To illustrate, from the external point of view one might ask about the content and function of moral concepts. One might ask whether and how moral properties are sewn into the fabric of reality. And one might ask whether and how we can have knowledge of moral properties, if indeed there are such properties. As with the internal moral perspective, these are questions that we can discuss together, for the external perspective is one that any moral agent can occupy in a reflective moment. They are also questions about which we can disagree. But such disagreements are not straightforwardly moral, for they are disagreements about morality, and they occur outside of moral discourse. At least, this is where they seem to occur.

The distinction between internal and external perspectives helps to shed light on a commitment mentioned in the previous chapter. Robust Realism, as I understand it, accepts the following:

External Metaphysics. We can assess the nature of moral properties from an external metaphysical standpoint, and the existence of such properties is to be defended or rejected primarily from that standpoint by appeal to external metaphysical considerations.

Robust Realists understand the Realist Commitments in light of External Metaphysics. They hold that these claims about moral truths and moral properties can be developed, and then either defended or rejected, from the external standpoint. This is basically what puts the ‘Robust’ in Robust Realism.

¹ For the ‘detachment’ claim, see Miller (2009).

2.2 Anti-Archimedean Challenges

Ronald Dworkin, a Relaxed Realist, describes the external metaphysical perspective as ‘Archimedean.’² According to Dworkin, theories that claim to occupy this external standpoint are Archimedean in that they “purport to stand outside a whole body of belief, and to judge it as a whole from premises or attitudes that owe nothing to it” (1996: 88). But philosophers like Dworkin deny that one can engage in such Archimedean metaethical theorising. This puts us in a position to understand the challenge posed by Quasi-Realism and Relaxed Realism. Both theories are, in a certain sense, *anti-Archimedean*. In one way or another, they deny that we can make sense of the Realist Commitments from an Archimedean point of view.

If the Quasi-Realist and the Relaxed Realist are right to reject Archimedean metaethics, and if they are also right to interpret External Metaphysics as an Archimedean commitment, then the very possibility of defending Robust Realism is undercut, for there is simply no way to even articulate its external metaphysical commitments. This is the sense in which Quasi-Realism and Relaxed Realism entail that there is no *space* for Robust Realism. And if there is no space for Robust Realism, the only way to retain the Realist Commitments is by treating them as moral. In other words, Quasi-Realists and Relaxed Realists think that we can only keep these commitments as internal moral claims. We cannot keep them if they are external metaphysical claims that require an Archimedean point of view.

This makes it clear that we should not think of anti-Archimedean as moral sceptics. Quasi-Realists and Relaxed Realists can and do accept the Realist Commitments. They just deny that these commitments involve any robust metaphysics. Indeed, they typically deny that they even *could* come with any robust metaphysics – given their anti-Archimedean views, Quasi-Realists and Relaxed Realists see an internal moral interpretation as the only available way of capturing the Realist Commitments. In short, the claim that there exists an irreducible, non-natural, mind-independent, and categorical moral reality simply *must* be seen as a claim made within ethics.

² Dworkin is not the first to use this term in relation to ethics, as Bloomfield (2009: 285) notes. For slightly different senses of ‘Archimedean’ in ethics, see Williams (1985) and Gauthier (1986).

So far, I have lumped the Quasi-Realist and the Relaxed Realist together. There is a rationale for such a lumping, for the two theories are closely connected – they are both forms of anti-Archimedeanism. But the picture is a little more complex than this suggests, for Quasi-Realists and Relaxed Realists are anti-Archimedean in importantly different ways, and these different styles of anti-Archimedeanism correspond to two distinct challenges facing the Robust Realist. So, before I argue that it is at least possible to provide external metaphysical interpretations of the Realist Commitments, I'll first distinguish and clarify the two anti-Archimedean threats.

Of the two anti-Archimedean challenges, one is more modest and one is more ambitious:

Ambitious Anti-Archimedeanism. It is impossible to articulate an external metaphysical interpretation of the Realist Commitments. This is because the Archimedean standpoint from which such an interpretation would have to be articulated either does not exist, or else cannot be occupied by creatures like us.

Modest Anti-Archimedeanism. It is impossible to articulate an external metaphysical interpretation of the Realist Commitments. This is because, although the Archimedean standpoint from which such an interpretation would have to be articulated does exist and can be occupied by creatures like us, there is no way to construe the Realist Commitments at this standpoint.

Quasi-Realism is a version of Modest Anti-Archimedeanism, and Relaxed Realism is a version of Ambitious Anti-Archimedeanism. I'll explore and respond to both of these challenges, but my focus will be on the ambitious one. This is because, as I'll explain, it presents the deeper threat to Robust Realism.

To clarify, Ambitious Anti-Archimedeanism proposes a wholesale rejection of metaethics. It claims that there is no space at all for external theorising about morality. On this view, there is simply no such thing as a genuinely external metaethical theory. That sounds dramatic, and indeed it is. Modest Anti-Archimedeanism allows that there is an Archimedean standpoint, and is to that extent less dramatic than its more ambitious counterpart. Even so, Modest Anti-Archimedeanism can seem a threat to Robust Realism. In the next two sections I will therefore explain why it is not.

2.3 Quasi-Realism

Quasi-Realism is a form of Modest Anti-Archimedeanism. To get a grip on the Quasi-Realist picture we can turn to its originator – Simon Blackburn. According to Blackburn, the Quasi-Realist project is the attempt “to earn, on the slender basis, the features of moral language ... which tempt people to realism” (1984: 171). This makes it appealing to see Quasi-Realism as a sort of ‘diet realism,’ for it purports to offer all of the tasty realist treats without any of the fattening metaphysics. However, it has often been said that Quasi-Realists want to have their cake and eat it too, so it may be that ‘diet realism’ is not such an apt label after all. I will shortly explain the Quasi-Realist position in detail by outlining how it interprets the Realist Commitments, but first it is worth commenting more broadly on its main goals.

Quasi-Realism has been understood in a variety of ways – there is more than one way to interpret its goals. For instance, on one influential conception, Quasi-Realism is characterised by a concern to imitate what Robust Realism says about moral reality.³ Another conception is that it is characterised by a concern to imitate some of what Robust Realism says, but only some of it.⁴ Yet another does not involve any imitation claim, and is simply characterised by a concern to capture certain parts of everyday moral thought and discourse without a robust metaphysics.⁵ Interpreting the Quasi-Realist project can therefore be a tricky business, and different versions of it may be more or less threatening to the possibility of Robust Realism. Whatever the best interpretation of Quasi-Realism, many things that influential Quasi-Realists have said do read as denying space for a robustly realistic metaethical theory. However, we can easily defuse the threat that Quasi-Realism appears to pose for the Robust Realist. To see this, we need only look at the extent to which its anti-Archimedeanism is modest.

³ This reading, which is (fairly or unfairly) associated with Blackburn, seems particularly vulnerable to what Dreier (2004) calls ‘the problem of creeping minimalism.’ It’s a very common interpretation, but it’s probably not the most charitable.

⁴ This is one way of interpreting Gibbard (2002; 2003; 2011). It is better placed to avoid creeping minimalism, and perhaps arguments like that of Street (2011).

⁵ This interpretation is clearest in Ridge (2014; 2015) and Dreier (2015), though it may be what others have had in mind too. It seems well-placed to avoid creeping minimalism, and is likely the most charitable reading of Quasi-Realism.

I'll come to that in the next section, but let's start by considering how Quasi-Realists see the Realist Commitments. First, to deliver moral truths the Quasi-Realist commits to a deflationary theory about truth in general.⁶ According to deflationary theories, predicates like '... is true' and 'it is true that...' can be eliminated from a sentence without any loss of meaning. The truth predicate does play important pragmatic roles – we use it to emphasise our convictions, and it enables speakers to endorse multiple propositions without having to explicitly assert each and every one of them in turn – but on a deflationary theory there is no semantic difference between '*p*' and 'it is true that *p*.' Once a deflationary theory of truth is accepted, the claim that there are moral truths can be had on the cheap.

To illustrate, deflationists will read 'it is true that killing is wrong' as semantically equivalent to 'killing is wrong.' And 'killing is wrong' is an internal moral claim, the meaning of which has yet to be elaborated. Now, the Quasi-Realist standardly defends 'Expressivism' about the meaning of moral sentences.⁷ According to Expressivists, the meaning of a sentence comes from the state(s) of mind that the sentence serves to express. Expressivists deny that there are robust moral properties, but they also deny that moral sentences aim to refer to such properties. They see moral sentences as serving to express non-representational states like approval or disapproval, and not representational states like belief. Expressivists thus see 'killing is wrong' as expressing (something like) disapproval of killing. Now, if we say that 'it is true that killing is wrong' is semantically equivalent to 'killing is wrong,' and if 'killing is wrong' is simply a moral claim that expresses disapproval of killing, then we can get moral truth without accepting anything metaphysically robust along the way. For the sentence 'it is true that killing is wrong' simply expresses disapproval of killing. So, adding the truth predicate to a moral sentence doesn't push us outside of moral discourse, and in that sense Truth is an internal moral claim.

⁶ Blackburn (1996: 86; 1998: 78-79). For a different view, see Ridge (2014: Ch. 7).

⁷ I say 'standardly' because I interpret Quasi-Realism and Expressivism as distinct and separable aspects of one theoretical package. This has the potential to confuse, for some treat 'Expressivism' and 'Quasi-Realism' as synonymous. I imagine that these people are a minority, for this use of terms unhelpfully masks important possibilities. Specifically, we should allow that one can see moral judgement as fundamentally non-representational without committing to Quasi-Realism, and we should also acknowledge that Expressivism is not the only view of moral discourse with which Quasi-Realism could be combined.

How about Properties? Well, according to Blackburn, there isn't any harm "in saying that ethical predicates refer to properties, when such properties are merely the semantic shadows of the fact that they function as predicates" (1993a: 181).⁸ And, if Quasi-Realists take moral properties to be 'semantic shadows,' as Blackburn puts it, they can also understand such properties as irreducible and non-natural. That our moral concepts are irreducible and non-natural is something that a lot of metaethicists now accept. This is because the 'Open Question Argument' defended by G.E. Moore (1903) can appear to have important semantic implications, even if its metaphysical implications have been overstated.⁹ By treating moral properties as 'semantic shadows,' the Quasi-Realist can then read Non-Reductivism and Non-Naturalism at the semantic level – that is, as located within moral discourse – rather than at the external metaphysical level.

One might think that it will be harder for Quasi-Realists to deliver Mind-Independence. Given that their theoretical package also includes Expressivism, they may appear to be committed to treating morality as dependent on the attitudes of agents (their approvals or disapprovals, or whatever). However, the Quasi-Realist also sees Mind-Independence as an internal moral claim. This might sound odd, and I will discuss it further in a moment, but here is the basic idea as Blackburn presents it:

'[M]oral truths are mind-independent' can *only* summarise a list like 'If there were no people (or people with different attitudes) then *X* ...' where the dots are filled in by some moral claim about *X* (1998: 311).

On this view, to say that a moral claim is mind-independently true is to make a counterfactual moral claim. It is to apply one's moral judgement to a range of counterfactual scenarios, including scenarios in which one's attitudes differ.¹⁰ Quasi-Realists thus see Mind-Independence as a moral claim.

⁸ I take it that Blackburn is gesturing here at the 'abundant' conception of properties. For more on 'abundant' and 'sparse' properties, see Lewis (1983; 1986) and Schaffer (2004).

⁹ From the openness of 'this is pleasurable, but is it good,' we can plausibly infer that 'good' doesn't mean the same as 'pleasurable.' But, contra Moore, it doesn't follow that goodness and pleasurableness are distinct properties, for this would require a synonymy constraint on property identity. And such a constraint is subject to well-known counterexamples (consider 'water' and 'H₂O'). Cf. Gibbard (2002). For a different view, see Cullison (2009).

¹⁰ Cf. Gibbard (2003: 186).

They are likely to make a similar move in attempting to capture the commitment to Categorical Authority. Remember that the authority of categorical direction is supposed to be independent of all social and psychological contingencies. Although this independence claim concerns the authority of moral direction, rather than the existential status of moral properties, the Quasi-Realist can again understand it simply as the application of an ordinary first-order moral judgement to a wider range of counterfactual scenarios. So, for the Quasi-Realist, to think that *A* has a categorical reason to Φ is just to think that *A* would have a reason to Φ even in worlds in which the social or psychological situation is different. To make a claim about the categorical authority of moral direction is just to have a certain sort of moral thought, requiring nothing metaphysically robust.¹¹

In sum, Quasi-Realists deny that we can have a metaphysical view of the Realist Commitments, but offer a moral reading instead. Of course, this has been a brief overview of Quasi-Realism. A great deal more could be said. However, I can now show that Quasi-Realism does not constitute a serious challenge to the *possibility* of developing a robust metaethical theory.

2.4 Foundations of Quasi-Realism

The modesty of the challenge posed by Quasi-Realism is the result of its reliance on certain Archimedean foundations. To see this, consider the Quasi-Realist's interpretation of Mind-Independence. Blackburn stresses that the *only* available interpretation of this commitment is a substantive moral interpretation:

The wrongness of wanton cruelty does indeed depend on things – features of it that remind us of how awful it is. But locating these is giving moral verdicts. Talk of dependency is moral talk *or nothing* (1993a: 173, my emphasis).

Claims about whether or not morality is mind-independent are made within ethics, on this view, and can *only* be made within ethics. There is therefore no way to make sense of an external metaphysical construal of Mind-Independence. In that sense, there is simply no space for a robustly realistic interpretation of this commitment. Or, indeed, Robust Realism in general.

¹¹ For detailed discussion, see Gibbard (1990: 162-165), Blackburn (1998: 266; 2010b), and Lenman (2013: 401-402).

However, shortly after suggesting that talk of morality's being or failing to be mind-independent is moral talk *or nothing*, Blackburn says the following:

[T]here *would* be an external reading if realism were true. For in that case there would be a fact, a state of affairs (the wrongness of cruelty) whose rise and fall and dependency on other things could be charted. But anti-realism acknowledges no such state of affairs and no such issue of dependency (1993a: 173).

This is Blackburn's account of what it would take for there to be a robust reading of Mind-Independence, and it clarifies that Quasi-Realists are only able to deny that there can be such a reading because they accept a metaphysically anti-realist theory – Expressivism – at the meta-level. In other words, Quasi-Realism is built on certain Archimedean foundations. This is not news. After all, Blackburn is explicit about it. However, it is important for seeing the modesty of the Quasi-Realist's anti-Archimedean threat.

Recall that the Modest Anti-Archimedean doesn't think that the impossibility of giving a robust construal of the Realist Commitments is due to the lack of an Archimedean metaethical standpoint. They therefore have to offer some alternative support for the claim that there can be no external metaphysical reading of, say, Mind-Independence. And what the Quasi-Realist offers is the following: *Expressivism is the true meta-level theory*. That is, the possibility of offering a robust view of moral reality is locked out by the truth of Expressivism. Put another way, Expressivism is the foundation on which a Quasi-Realist's anti-Archimedean challenge is built.

Again, this is not a new observation – it's built into the theoretical package itself. The significance of making it explicit is that it exposes an important dialectical point, which will in turn clarify how Quasi-Realism fails to constitute a serious challenge to the possibility of articulating and defending a robust metaethical theory. The dialectical point is as follows: in defending themselves against the Quasi-Realist's anti-Archimedean challenge, the Robust Realist can simply give their reasons for rejecting Expressivism. They can work on the basis that there is space for Robust Realism because they take that Archimedean theory to be false. And, if Expressivism is false, then there is no foundation for the Quasi-Realist's claim that Mind-Independence cannot be read as an external metaethical view.

Put another way, the real challenge to Robust Realism comes from the Expressivist part of the package rather than the Quasi-Realist part. The real challenge is Archimedean rather than anti-Archimedean, for it is this Archimedean theory that does the heavy lifting against an external metaphysical reading of Mind-Independence. Indeed, the Quasi-Realist's modest anti-Archimedean challenge would collapse if it were to lose its Expressivist foundation. In short, the important moment of disagreement between Robust Realists and Quasi-Realists occurs at the external level, and Robust Realists should engage the Quasi-Realist at this foundational stage.

And this is just to say that Robust Realists can defend themselves against Quasi-Realism by doing something that they've always done, that is, by rejecting Expressivism. Quasi-Realists won't be able to convince you that there is no space for Robust Realism unless you've already been sold on Expressivism at the meta-level. And this is something on which Robust Realists remain unsold. So – and here's the crucial point – Robust Realists do not need any *new* strategic manoeuvres to answer the Quasi-Realist challenge. They can just continue making their arguments against the Expressivist part of the package.¹² More broadly, if an anti-Archimedean challenge can only get off the ground by appealing to a certain external metaethical theory, it is not a serious threat to the *possibility* of Robust Realism.

I imagine that some Quasi-Realists will say that they are fine with this conclusion. The extent to which Quasi-Realists view themselves as challenging the *possibility* of Robust Realism will depend significantly on how they interpret the goals of their project. And, as I've noted, different interpretations are available. So it may be that, despite the rhetoric of certain prominent Quasi-Realists, their project does not have to be seen as an attempt to paint Robust Realism out of the picture. Fine by me. My current concern is just to show that it is possible to articulate a robust moral metaphysics. If there are Quasi-Realists who never deny that this is a possibility, then I have no beef with them. My topic has been a more intolerant-sounding Quasi-Realism, and I have shown how this view fails to present a threat to the possibility of developing a robust metaethical theory.

¹² There are many things that the Robust Realist might say in rejecting the Expressivist view. Some interesting critiques in the recent literature are developed by Unwin (1999), Cuneo (2006), Schroeder (2008), and Woods (2014).

I should make clear that I have *not* shown, and have not attempted to show, that the Expressivistic Quasi-Realist package is false. For all I've said so far, it may be that we should in the end accept it. I have only been arguing that it fails to constitute a genuine challenge to the *possibility* of Robust Realism. By accepting a point of view from which one can argue about external metaphysical interpretations of the Realist Commitments, even the most intolerant forms of Quasi-Realism make space for Robust Realists to articulate and defend their external metaphysical position. Robust Realism may still turn out to be false, but there is room for such a theory.

Here is the lesson. An anti-Archimedean threat that builds on an Archimedean foundation does not constitute a genuine challenge to the *possibility* of our articulating a robust moral metaphysics, for one can respond to such a threat simply by engaging it at its foundational level. The deeper anti-Archimedean threat will come from theories that reject any such level at all, and thus refuse to be engaged at it. In short, the more serious challenge comes from the Ambitious Anti-Archimedean theories. Modest Anti-Archimedeanism doesn't require new strategic manoeuvres from the Robust Realist, but Ambitious Anti-Archimedeanism does. I will soon provide the Robust Realist some new manoeuvres to deal with the ambitious challenge. First, however, I will explain the challenge in greater detail.

2.5 Relaxed Realism

Ambitious Anti-Archimedeanism is brought to life by Relaxed Realism. The position I have in mind is often called 'Quietism,' but its defenders tend to distance themselves from that term. They understand themselves as defending the only feasible sort of realism, and claim that efforts to be more robust are confused. My preferred label – 'Relaxed Realism' – is apt, for philosophers who defend this view combine the Realist Commitments with what Sarah McGrath describes as "a certain lack of anxiety about the status and standing of morality, despite understanding morality in ways that might naturally encourage such anxiety" (2014: 187). As will become clear, this lack of anxiety is due to their Ambitious Anti-Archimedeanism. Unlike the Quasi-Realist, the Relaxed Realist doesn't rely on any external metaethical foundation with which the Robust Realist can engage, and this is what makes them a much deeper threat to the possibility of Robust Realism.

For the Relaxed Realist, the Realist Commitments are just internal moral claims. *And that's it.*¹³ There's no more to say. Attempts to give an account of the metaphysical status of those commitments are misguided, as are attempts to explain what we are doing when we talk about such commitments. They are moral rather than metaphysical claims, and there is nothing else to it. This might sound puzzling, but there are advantages to a Relaxed Realist approach that make it attractive.¹⁴ After all, if Relaxed Realists are correct, any external metaphysical doubts that you may have about fitting morality into a scientific worldview are misplaced. After all, if there is no external standpoint from which wholesale moral scepticism can make sense, then one can only doubt the Realist Commitments at the internal perspective, and these doubts must come from moral rather than metaphysical factors. Relaxed Realism thereby seems to undercut certain sceptical views about moral reality, for it says that one can only question a Realist Commitment by making a moral case against it. Its ability to deal with a major form of scepticism about morality is an important attraction of Relaxed Realism. So, to clarify how this position is supposed to work, let's see how Relaxed Realists interpret the Realist Commitments in more detail.

Relaxed Realists typically suggest that there is an independent or strictly 'autonomous' domain of normative discourse. Just as there are autonomous domains of mathematical and scientific discourse, there is a group of normative concepts and claims that constitutes the autonomous normative domain. And, for Relaxed Realists, the correct view of Truth sticks within this normative domain. As Dworkin puts it, it is a mistake to "expect answers that step outside morality to find a nonmoral account of moral truth ... that expectation is confused: it rests on a failure to grasp the independence of morality" (2011: 38). More specifically, T.M. Scanlon says that "truth values of statements about one domain, insofar as they do not conflict with the statements of some other domain, are properly settled by the standards of the domain they are about" (2014: 19). As long as normative statements don't conflict with (say) scientific statements, their truth is to be decided by standards internal to normative domain. And these are just normative standards. That is, substantive principles or norms.

¹³ More strictly, some Relaxed Realists – like Scanlon – are constructivists about the moral but relaxed about the normative. This is a wrinkle that we can set aside for now.

¹⁴ Cf. McPherson (2011: 238).

Relaxed Realists also claim that there are moral properties in a sense that require no second-order metaphysics. Derek Parfit has called himself a ‘non-metaphysical cognitivist,’ for example, but he allows that there can be normative properties in a minimal sense:

I use the word ‘property’ in the wide non-metaphysical sense with which we can restate any claim that is, or might be, true. Whenever someone ought to act in some way, for example, we could say either that this act has the property of being what this person ought to do, or that this person has the property of being someone who ought to act in this way (2011 v.2: 756).

I don’t know if I understand Parfit’s overall view, but he is clearly relaxed about normative properties.¹⁵ He doesn’t view them as metaphysically robust.

Given their appeal to an autonomous normative domain, it’s easy to see how Relaxed Realists view Non-Reductivism and Non-Naturalism. Normative predicates belong to an autonomous domain of discourse, so they won’t reduce to predicates from other (descriptive) domains. Given the irreducibility of normative predicates, the properties picked out by such predicates will be irreducible too. In sum, given the autonomy claim, the Relaxed Realist can suggest that moral and descriptive properties are distinct.

Relatedly, the Relaxed Realist needn’t think that normative truths are discoverable by methods appropriate for, say, the scientific domain. Truths about normative properties will be beyond the purview of the sciences, for they are part of a distinct domain and are thus known by methods appropriate for that domain. Perhaps a form of foundationalism is correct, where the fundamental normative truths are detected through direct intuition. Or perhaps we should adopt a coherentist method, where we seek a reflective equilibrium among considered judgements.¹⁶ Either way, Relaxed Realists understand Non-Reductivism and Non-Naturalism such that they stay at the moral level, avoiding anything metaphysically robust.

¹⁵ Cowie (2014a) helpfully seeks to interpret Parfit and other Relaxed Realists. Note that Parfit, in the forthcoming third volume of *On What Matters*, changes his preferred label to ‘non-realist cognitivism.’ I take it that his view still counts as Relaxed Realist in my sense.

¹⁶ For the former approach, see Parfit (2011 v.2: Ch. 32). For the latter, see Dworkin (2011: 82-86) and Scanlon (2014: Ch. 4).

When it comes to Mind-Independence, the Relaxed Realist once again understands this as a commitment that can only be articulated and defended or rejected at the internal moral level. As Thomas Nagel puts it, it is only by “thinking about what to do and how to live ... [that] we can find methods, reasons, and principles whose validity does not have to be subjectively or relativistically qualified” (1997: 102). This is how Relaxed Realists offer a non-metaphysical construal of Mind-Independence. They say that it only makes sense as a first-order moral commitment, one that must be defended by first-order moral argument. Because thinking about what to do and how to live is simply internal moral thinking, not external metaphysical thinking, it doesn’t commit us to anything metaphysically robust.

Moreover, if the normative domain is autonomous, the authority of normativity will not depend on social and psychological contingencies lying outside of that domain. We can see this by considering the notion of a reason. A reason to Φ is a consideration that counts in favour of Φ -ing. According to many Relaxed Realists, the concept of a reason is basic. That is, despite what some say, no further explanation of it can be given:

[A] further explanation might be an explanation of what “counting in favour of” amounts to. This might take the form of an explanation of the “grip” or “authority” of reasons, of the kind offered by Kantians and others ... however, it seems to me that no such further explanation of reasons need or can be given: the “grip” that a consideration that is a reason has on a person for whom it is a reason is just being a reason for him or her (Scanlon 2014: 44).

It is misguided, in other words, to seek some deeper explanation for the authority of reasons. The concept of a reason is fundamental. An agent’s having an authoritative reason to Φ thus requires no further explanation in terms of its being conditional on a desire, convention, commitment, or agentive status. It is primitive, and does not take us beyond the normative domain.

In sum, Relaxed Realists reject a metaphysical view of the Realist Commitments, claiming that they can be located within the autonomous normative domain. Moreover, they do not rest this internal reading of the Realist Commitments on any Archimedean foundation. Their challenge is thus ambitious, constituting a serious threat to the possibility of Robust Realism.

2.6 Compatibility

The threat posed by Relaxed Realism cannot be defused by engaging it at its external metaethical foundation, for Relaxed Realists claim that there is no space for external metaethics. This is the nature of their Ambitious Anti-Archimedeanism. But how can it actually be shown that there is no space for external metaethics? We've seen internal moral readings of the Realist Commitments, but why should we accept that these are the *only* available readings? This question is crucial, for if an external metaethical reading is also available, the more relaxed option might appear hollow by comparison. In considering this matter I will use Dworkin's work as my jumping off point. McGrath rightly says that Dworkin's defence of Relaxed Realism is "undiluted and uncompromising" and thus "provides a useful case study" for critics to examine (2014: 187). Dworkin's suggestions reveal a lot about the Ambitious Anti-Archimedean challenge to Robust Realism.

In attempting to show that an internal moral view of the Realist Commitments is the only one available, Dworkin (1996: 96-97) proposes a test to help us assess the viability of Archimedean metaethics. The test starts with a moral judgement. Consider, for example, my judgement that killing is wrong. This is clearly a moral judgement. We can locate it within the internal point of view. But imagine that I decide to elaborate on it by saying that it is *true* that killing is wrong, and that wrongness is a *property* of killing. I may also say that this property is *irreducible, non-natural, and mind-independent*. I may add that it involves some *categorical direction*. To stress my conviction, I may even say that killing is *really and actually* wrong.

Dworkin brands these the 'Further Claims,' and he suggests two questions that one must ask about them in order to test the viability of Archimedean metaethics. First, is there a way of interpreting the Further Claims as moral statements? That is, can the Further Claims be viewed as restatements or clarifications of the original moral judgement? Second, is there a plausible way of interpreting any of the Further Claims as morally neutral? Dworkin tells us that, if the answer to the first question is *yes* and if the answer to the second question is *no*, we cannot do Archimedean metaethics. And, given that the Robust Realist's commitment to External Metaphysics is apparently a commitment to Archimedeanism, this would mean that there is no external metaphysical perspective from which one can even state that theory. This would undercut the possibility of Robust Realism.

According to Dworkin, the answer to the first question is indeed *yes* and the answer to the second question is indeed *no*. That is, there is a way to read the Further Claims as internal to ethics, and there is no way to read them as morally neutral. Archimedeanism is thus undermined. So, if we are assuming that External Metaphysics is itself an Archimedean commitment, there is no space for a robust interpretation of the Realist Commitments. This is clearly a threat to my project, for my aim is to show how the Realist Commitments can and should be developed by the Robust Realist.

Why should one think that the answer to the first question is *yes*? Dworkin (1996: 97) suggests that the most natural reading of the Further Claims is an internal moral reading. When someone says that a certain moral claim is true independently of anyone's responses or attitudes, for example, the most natural interpretation (according to Dworkin) is that the individual in question is just clarifying the content of their substantive moral opinion. As it happens, I don't find that the most natural reading of this claim. Dworkin does, however, and the popularity of Relaxed Realism suggests that he is not alone. Anyway, whether or not it is the natural reading, I accept that the further claims can be read as moral claims. That is, I accept that the answer to the first question is *yes*. There is a way of reading the Further Claims such that they fall within the internal point of view.

But why is the answer to the second question *no*? Well, according to Dworkin, claims that are seen by philosophers as metaethical actually end up being or entailing substantive moral claims:

The philosophical-sounding proposition that there are moral properties in the universe, for example, is or entails ... that some acts really are unjust, or some people really are good, or something of the sort. So read ... a skeptic who denied it would hardly be neutral toward substantive morality (1996: 100).

In short, if you accept that moral properties exist, you must accept that they are instantiated in certain acts and people (and such). But this means that even abstract claims like 'there are moral properties in the universe' are not morally neutral, for they entail substantive moral propositions. Now, if such claims are not morally neutral, they can't be Archimedean. And, for Dworkin, if these claims aren't Archimedean they aren't external and metaphysically robust. The very possibility of Robust Realism is thus undercut.

This leaves Relaxed Realism as the only available interpretation of the Realist Commitments. And without a more robust reading of those commitments, a relaxed reading may seem less hollow. Should my project in subsequent chapters thus be seen as an account of the internal moral commitments of Relaxed Realism, rather than the external metaphysical commitments of Robust Realism? I think not. Against what the Ambitious Anti-Archimedean suggests, external and metaphysical commitments are entirely possible. I will show that Dworkin overstates the case for being relaxed instead of robust, arguing that his mistake is to think that external metaphysical commitments require an Archimedean point of view. In this way, the Robust Realist can deal with the Ambitious Anti-Archimedean threat.

In showing exactly how Robust Realists can handle this threat, it will be productive to consider the following thesis:

Compatibility. External metaethical theories and claims must be compatible with every internal moral theory and claim.

Compatibility is to be understood as a constraint on which theories can count as metaethical.¹⁷ It explicitly ties the external to the Archimedean. According to Compatibility, a theory that fails to be compatible with every first-order moral view thereby fails to count as external and metaethical. Dworkin's test aims to show that no theory conforms to Compatibility. Even very abstract claims end up entailing substantive ethical positions, and supposedly external metaethical theories thus fail to be compatible with every internal ethical claim. In other words, they fail to conform to the constraint provided by Compatibility. So, if Compatibility genuinely constrains what can count as an external metaethical theory, there are no external metaethical theories. There is, in that case, no way to be a Robust Realist.

Compatibility, as I will explain, is not something that we have to accept as a constraint on metaethical discourse. It is open to the Robust Realist (and others) to reject Compatibility, and in doing so undercut the Ambitious Anti-Archimedean threat. Having said that, it is worth noting that Compatibility has lurked in the background of many metaethical views.

¹⁷ McPherson (2008: 3) discusses a similar sort of constraint, which he names 'Neutrality.' It is better to put the point in terms of compatibility, however, for there are different ways of understanding moral neutrality that can figure in this debate. I discuss different forms of neutrality in §2.7.

For instance, the early 'Emotivists' – who were ancestors of the contemporary Expressivists – were certainly sympathetic to the idea that to do metaethics is to abstract away from first-order moral questions. A.J. Ayer suggests that strictly philosophical work on ethics should “make no ethical pronouncements” (1936: 105). Charles Stevenson aims to “retain that difficult detachment which *studies* ethical judgements without *making* them” (1944: 110). P.H. Nowell-Smith neatly captures this view of substantive moral theorising when he says that “[a] philosopher is not a parish priest or Universal Aunt or Citizens’ Advice Bureau” (1954: 12). I am not entirely unsympathetic to this line of thought, for I can attest that philosophy PhD programmes are not production lines for moral saints.¹⁸ Anyway, these thinkers were sympathetic to Compatibility, or something like it.

And they are not alone. J.L. Mackie – an Error Theorist, rather than an Emotivist or an Expressivist – also bought into something like the Compatibility constraint. Mackie saw first-order ethics as distinct from second-order metaethics, such that “one could be a second-order sceptic without being a first-order one, or again the other way round” (1977: 16). In short, it is clear that Compatibility has had an important influence over some philosophers. Dworkin, however, thinks that no theory conforms to Compatibility, and he takes this to undermine Archimedean metaethics. This, in turn, is meant to undermine the very possibility of engaging in external metaphysical theorising. Robust Realism therefore appears to be in some trouble. After all, if there is no point of view from which one can articulate (let alone defend) the external metaphysical commitments of a theory, that’s more than a little inconvenient for those who understand themselves as subscribing to a robustly realistic account of moral reality. This is the essential core of the Ambitious Anti-Archimedean challenge as Dworkin develops it. This relaxed approach to morality seems to pull the Archimedean rug from under the Robust Realist’s external metaphysical feet.

¹⁸ Less flippantly, whilst a view of philosophy that treats moral theory as unphilosophical is obviously too narrow, there is something to be said for the idea that one should not expect to derive specific moral verdicts from highly abstract claims made in metaethics. Not everyone would agree with this – Hare (1981) is the obvious counterexample – but it seems plausible to me. One might worry that rejecting Compatibility, as I do below, fails to respect this plausible idea. I respond to this worry in §2.7 by arguing that there is a more modest version of this constraint that does not have this implausible result.

2.7 Against Compatibility

How can Robust Realists answer this challenge? Some interesting replies to Dworkin (and to Relaxed Realist positions more generally) try to show that it is in fact possible to make external metaethical claims that conform to Compatibility.¹⁹ In other words, some responses seek to vindicate the Archimedean standpoint. However, this is not the sort of response that I am going to offer. Instead, as I have already indicated, I will suggest that Robust Realists can reject Compatibility without thereby losing the ability to articulate and defend a robust construal of the Realist Commitments. After all, it's not clear why we are meant to suppose that Robust Realists are committed to Compatibility. If they were to reject this constraint on what can be counted as an external metaethical theory, will the heavens (I want to say Plato's heavens) fall? I doubt it. In particular, we can take issue with the suggestion that an apparently metaethical claim's merely *entailing* a substantive moral claim somehow shows that it was moral all along.

I explained in §2.6 that, according to Dworkin, a supposedly metaethical claim's having substantive moral bearings ensures that it is itself a substantive moral claim. For instance, the abstract claim that there are moral properties in the universe might sound as though it's entirely non-committal on first-order matters, but it entails that "some acts really are unjust, or some people really are good, or something of the sort" (Dworkin 1996: 100). It is therefore *not* an external metaphysical claim, according to Dworkin, for it fails to conform to Compatibility. It fails the Compatibility test, and thus fails to be Archimedean, and thus fails to be robust. But it is not easy to make sense of this move. How does some theory's having a substantive moral bearing make it a substantive moral theory? What would the general principle be? Perhaps that a claim that seems to belong to claim-kind₁ actually turns out to belong to claim-kind₂ if it entails a claim that belongs to claim-kind₂. But this would be a rather peculiar principle, as one can see by looking at some examples from other contexts.

¹⁹ The availability of Archimedean theories is defended by Dreier (2002) and Ehrenberg (2008). Note that, whilst my reply to the Ambitious Anti-Archimedean shows that Robust Realists need not rely on Compatibility, this does not necessarily mean that *no* metaethical theory is compatible with everything at the moral level. Just that it doesn't have to be to count as metaethical. So, even though we should not tie the external to the Archimedean, it may be that some theories do count as Archimedean.

Consider a theological claim – *God created the universe in seven days* – that implies a certain metaphysical claim – *naturalism is false*. Does the theological claim thereby fail to count as theological? It seems not. Is theology just part of metaphysical discourse? I don't see it. Or consider a neurological claim – *c-fibres are firing* – that implies a phenomenological claim – *pain is felt*. Does the neurological claim thereby fail to count as neurological? Does it just become phenomenological? Again, it seems not. And it would be peculiar to think that it did. In short, it is not appealing to suppose that claims belonging to claim-kind₁ actually turn out to belong to claim-kind₂ given entailments between the former and the latter. If this is what Dworkin has in mind, there is no reason to think that his claims succeed in undermining the possibility of providing a robust metaethical theory.

However, there is a better way to understand the point about entailments. The examples just given – the theological and neurological entailments – are (I presume) not instances of strictly logical entailment. And this is the sort of entailment that Dworkin and other Relaxed Realists are interested in.²⁰ This is not always clear from their writings, but maybe the suggestion is that logical entailments from the putatively metaethical to the substantively ethical show that there isn't any external metaethical level.

But the example that Dworkin himself provides is only a case of logical entailment when interpreted charitably. Recall that he offered the following:

The philosophical-sounding proposition that there are moral properties in the universe, for example, is or entails ... that some acts really are unjust, or some people really are good, or something of the sort (1996: 100).

The phrase “or something of the sort” is obviously crucial here. Strictly speaking, the putatively metaethical claim – *there are moral properties in the universe* – only logically entails something like – *certain aspects of the universe that can bear moral properties do in fact bear moral properties*. And why should Robust Realists be unsettled by that logical entailment? After all, it sounds very much like a metaphysically committed position. It thus seems as though there is no clear reason to suppose that the *logical* entailment undermines one's commitment to something *metaphysically* robust.

²⁰ Thanks to Matthew Kramer for suggesting this reading of Dworkin.

Put another way, Dworkin's example provides us with reason to reject Compatibility as a constraint on external metaethics. It does not provide us with a reason to reject external metaethics itself. It is therefore very tempting to allow that an external metaethical claim can fail to be compatible with every substantive moral theory and claim, even though it is a metaphysically committed claim *about* ethics. Robust Realists can thus turn the tables on Dworkin, pulling the anti-Archimedean rug from under his anti-metaphysical feet by saying that a Compatibility constraint is not something to which they were ever in fact committed. In sum, they should agree that Robust Realism fails to conform to Compatibility whilst denying that this has anything like the significance that Dworkin has to assume.

To illustrate, look at how Moore's answers to certain metaethical questions seem to bear on his acceptance of a certain first-order moral position. Specifically, his commitment to some form of utilitarianism is implied by his definition of 'right' in terms of 'good.' Moore thought that the term 'good' was itself indefinable (and thus contrasts himself with 'analytic utilitarians' like Jeremy Bentham), but he held that to say that an act is right is to say that it promotes the good: whatever things are good, the right thing to do is to promote them.²¹ Moore's question about the meaning of 'right' is naturally seen as metaethical, the answer to which entails a (fairly abstract) claim within first-order ethics. My view is that such examples cast doubt on Compatibility as a constraint on metaethics, rather than on metaethics itself. A view can count as metaethical, and can involve metaphysical commitments, whilst also having implications for ethics.

Of course, in rejecting the Compatibility constraint we do not lose neutrality altogether. Even the Relaxed Realist agrees that a metaethical claim's entailing some ethical claim does not thereby mean that it cannot be compatible with other ethical claims. For instance, even if the claim that moral properties are sewn into the fabric of reality implies that some acts are unjust or that some people are good or something like that, it does not entail anything about *which* acts are unjust or *which* people are good. Nor does it tell us anything about what *makes* some acts unjust or some people good. So even if a metaethical theory fails to be compatible with all first-order moral claims, it can be compatible with very many of them.

²¹ For details, see Moore (1903: §14).

In particular, it can be compatible with the verdicts that we arrive at for highly specific moral questions. So, to reject Compatibility is not to reject first-order neutrality altogether. We can reject Compatibility and instead accept something like the following:

Modest Compatibility. Most external metaethical claims are compatible with a wide range of internal moral claims, including most claims about which verdicts are correct in particular cases.

I mention this because it preserves what is correct about Nowell-Smith's observation that philosophers (or metaethicists, at least) are not parish priests. Modest Compatibility accommodates the fact that specific moral guidance is unlikely to fall out of philosophical metaethics of an abstract sort.

Not everyone who rejects Compatibility will also endorse Modest Compatibility. For instance, some philosophers have thought that quite specific moral views can be derived just from the logical features of moral words.²² I believe that we should endorse it, however. After all, you'd be unlikely to consult a specialist in metaethics to find out whether eating meat is wrong. (At least, if you were to do so, it is unlikely that you would be appealing to them in their capacity as a metaethicist.) Moreover, one can be morally well-informed whilst knowing nothing of metaethics, just as one can be morally ill-informed whilst knowing a lot of metaethics, and it is hard to see how metaethical considerations fix a verdict on specific moral debates about euthanasia, meat eating, abortion, the environment, charity, war, and so on. Modest Compatibility manages to capture the fact that, whilst metaethics has *some* practical relevance, this will only go so far.

Before moving on, it is also worth noting that there are different ways of understanding neutrality.²³ Compatibility – which seems to be the most common interpretation of metaethical neutrality among those who actually articulate their view of it – fleshes out a metaethical theory's neutrality as consisting in its being compatible with every first-order theory and claim. But there is at least one form of neutrality that has nothing to do with Compatibility at all. Neutrality can be an *epistemic* matter.

²² Cf. Hare (1981).

²³ Cf. Gewirth (1968).

Even if an answer to a metaethical question entails an answer to a substantive question, it may sometimes be possible to ask and answer that metaethical question from an *epistemically* neutral standpoint. Doing so will involve suspending judgement on the answer to the substantive question. To illustrate, consider the following case. *Question*: how can I have moral knowledge? *Answer*: I am justified in believing that all of my moral judgements are true because I have a perfectly reliable faculty of moral intuition. Ignore the fact that this answer is both silly and hubristic. What matters is that both the question and the answer are metaethical. But the answer has substantive implications, for it entails that all of my first-order moral judgements are correct. For instance, it will entail that eating meat is wrong, assuming that I judge that eating meat is wrong. Still, whilst asking the question about moral knowledge, I can suspend judgement on the wrongness of eating meat. I can withhold assent on that moral question, and I can avoid even thinking about any substantive claims whilst comparing the merits of my hubristic intuitionism with its rivals.

My answer to the metaethical question about moral knowledge may not be logically compatible with every first-order theory and claim, then, but in arriving at that answer I am in some sense neutral. So, this is a way of being neutral about substantive moral matters. Of course, I am not suggesting that this sort of neutrality is Archimedean, or even that it will be possible in every case of external metaethical reasoning. However, it does indicate another way in which one can sometimes be neutral about moral claims even when those claims are entailed by one's metaethical views.

I don't know how often this can be done, and I'm not sure how easy it will be to achieve, but it may help to satisfy someone who remains disturbed by the idea that metaethics has even *some* practical import. I have encountered philosophers who are uncomfortable with the idea that metaethicists cannot abstract away from substantive ethical questions altogether. I find it hard to feel the full force of this concern, given that the substantive implications of most metaethical views will be fairly modest. But for those who do feel the concern, the epistemic neutrality to which I have pointed might provide some relief. If one feels unqualified to engage with substantive matters, or if one thinks that lack of neutrality on them gets in the way of good metaethics, it may be that (at least sometimes) we can suspend judgement on the ethical issues to look solely at metaethical issues.

So, by rejecting Compatibility, the Robust Realist can deal with the threat posed by Ambitious Anti-Archimedeanism. There is no reason to think that Robust Realists are committed to this constraint on metaethics, and rejecting it will not put them in danger of becoming parish priests if they accept Modest Compatibility instead. In short, the commitment to External Metaphysics can survive the Relaxed Realist challenge. We don't need to view it as an Archimedean commitment, not after we reject the Compatibility constraint. So, despite what Dworkin and others have said, there is space to articulate an external and metaphysical interpretation of the Realist Commitments. In other words, there is space to be a Robust Realist.

2.8 Moral and Metaphysical Argument

I have focused so far on Dworkin's uncompromising attack on metaethics. However, not all Relaxed Realists see themselves as attacking metaethics itself. For instance, Matthew Kramer is a Relaxed Realist who thinks that metaethical theorising is a thing, but who also thinks that the class of metaethical claims is located within the class of substantive moral claims. In other words, whereas I see the rejection of the Compatibility constraint as showing that robust external metaphysical commitments are possible despite logical entailments between metaethical and ethical propositions, Kramer sees this as revealing that metaethics is located within ethics itself. In short, whilst Kramer does not see metaethics as being external to ethics, he does see it as a second-order discipline about ethics. In his view, the entailments from the metaethical to the ethical reveal that this is a second-order discipline that is located within the first-order ethical domain.

Again, I am not sure that I understand this move. To show that a metaphysical claim about moral reality has logical implications for what we say in ethics is not to thereby show that the metaphysical claim is part of substantive ethical discourse. It seems to me a mistake to conflate these ideas. Perhaps this is terminological.²⁴ Perhaps some are willing to use terms like 'first-order' and 'moral' more expansively than others. If so, the dispute between Robust Realism and Kramer's form of Relaxed Realism is smaller than it looks. This would be a nice result. However, even if this is so, there is still a major difference between Kramer's view and Robust Realism.

²⁴ Cf. Enoch (2011: 130).

This difference lies in how Robust Realists expect arguments for the Realist Commitments to unfold. Owing to his view that metaethical discourse is just a part of moral discourse, Kramer thinks that the Realist Commitments are to be defended or rejected primarily by an appeal to substantive moral considerations. For instance, he maintains that certain theories – what he calls ‘Subjectivism’ and ‘Relativism,’ in particular – are forced to take substantive moral stands that expose them as untenable. For instance, Kramer tells us that, by making the correct moral principles constitutively dependent on our attitudes, the Subjectivist must accept that claims like ‘it is not the case that torturing babies for pleasure is morally wrong’ would be true in (for example) a world that contained no people.²⁵

Kramer thinks that this claim is obviously false. And, because he sees Subjectivist metaethical theories as being committed to such claims, Kramer maintains that they can be seen to be false as well. Put another way, Subjectivist metaethical theories are rendered unacceptable by their repugnant entailments at the substantive moral level. This point applies more generally. Kramer holds that metaethical positions are to be tested primarily by consideration of their substantive moral implications, and he judges that the Realist Commitments come out well from this sort of test.

In accepting External Metaphysics, a Robust Realist should deny this key role to moral argument. They should instead say that the Realist Commitments are to be defended or rejected primarily by an appeal to external considerations. We abstract away from substantive entailments and assess the metaphysical merits of various views of moral reality. This echoes the approach to metaphysical inquiry discussed in §1.2. On that approach, we develop various theoretical accounts of how reality might be. We then compare these theories by the standard criteria for theory choice, and provide arguments for thinking that reality is one way or another. In metaethics, we are of course concerned mainly with the moral aspects of reality. Robust Realists should thus say that a defender of the Realist Commitments needs to examine the competing theories of moral reality in order to assess their theoretical merits. As a way of arguing about moral reality, this is more attractive than the appeal to first-order positions.

²⁵ Kramer (2009: 30-35). I take it that a sophisticated subjectivist will in fact attempt to rigidify the subjective moral norms, but we can set this aside for the present.

To see how Kramer's way of arguing about moral reality leads to problems, consider that he frequently relies on the emphatic assertion of his personal moral convictions when attempting to establish the Realist Commitments, condemning certain Relativist theories, for example, by saying that they give us no way to morally criticise "fanatical Nazis and Maoists and other arrant villains" (2009: 45). Other metaethical views are likewise rejected because Kramer considers them "repellent" (2009: 32), or to have a certain "perniciousness" (2009: 33), or to be "appalling and crazy" (2009: 133). I generally find myself agreeing with Kramer's ethical judgements. I am no fan of fanatical Nazis and Maoists, and I probably wouldn't get on with most other arrant villains. However, it is unwise to give this sort of moral consideration primary significance in metaethical debate.

Note first that, from the claim that metaethics is part of ethics, it does not follow that ethically assessing metaethical positions has to be the only or even the primary way of arguing about moral reality. More crucially, focusing our attention on substantive convictions is unlikely to lead to progress in metaethics, however firmly they are held.²⁶ It is open to anti-realists to reply to Kramer by saying that what their arguments show is that, perhaps surprisingly, many substantive moral claims that look highly evident to us are in fact false (unless relativised, or whatever). Such a move seems to be dialectically acceptable, for if there is no strict division between ethics and metaethics then arguing from metaethical to ethical positions is just as feasible as arguing from ethical to metaethical views.

Assuming that the moral views to which Kramer appeals are not maximally evident, as one no doubt should, it is hard to see how Relaxed Realists can respond to the anti-realist without simply reasserting their personal convictions even more emphatically. But if that's the only move available to us, or even if it's meant to be our primary move, then progress seems unlikely. After all, this way of arguing just won't be convincing to those who are already prepared to reject certain Realist Commitments in light of metaphysical problems that they take to be associated with them. And that's significant, at least if we were hoping that our arguments might be dialectically effective. We are better off at the external metaphysical level, which turns on modest speculation about ways moral reality might be.

²⁶ I return to this sort of point in Chapter 6.

So, there are important differences between Robust Realism and Kramer's brand of Relaxed Realism. Specifically, they take different views about how metaethical arguments should unfold. And the Robust Realist has the less problematic approach here. I return to these dialectical issues in Chapter 6, for whilst I am confident that the Robust Realist's approach is less problematic than the Relaxed Realist's, there are (I will argue) far more general problems with how we argue about moral reality, problems that should be troubling to anyone with the optimistic goal of preaching beyond the choir. Anyway, what matters here is that Robust Realism is clearly not the same as Kramer's version of Relaxed Realism, and that on balance a robust view seems to have more going for it than his relaxed view.

2.9 Conclusion

There is space for Robust Realism. Despite what a number of both Modest and Ambitious Anti-Archimedean have claimed, it is entirely possible to develop and defend a metaethical theory that has external metaphysical commitments. Of course, none of this is to show that Robust Realism is actually the true metaethical theory. For all I have said, it may be that the Realist Commitments are all false. Maybe there are no moral truths, or no moral properties. Or maybe there are such truths and properties, but they are all reducible, natural, mind-dependent, and conditional on contingent desires or social conventions. However, we can at least say that there is room to *articulate* a robustly metaphysical view of these commitments. This may sound like a modest conclusion, but it matters a great deal given the apparent significance of the threat posed by anti-Archimedean views like Quasi-Realism and Relaxed Realism. Having defended the possibility of developing a robustly realistic metaethical theory, the next task will be to actually develop Robust Realism in detail. I begin that task in the next chapter.²⁷

²⁷ A shorter version of this chapter is published as Ingram (forthcoming).

3 The Supervenience of Moral Properties

Robust Realists think that there are substantive moral truths, and that at least some of these pertain to the moral properties of acts, or types of act. I will not elaborate on these claims, however, for in themselves they do little to distinguish Robust Realism from rival theories. Instead, it will be more productive to investigate the four Realist Commitments that make Robust Realism a uniquely uncompromising position: Non-Reductivism, Non-Naturalism, Mind-Independence, and Categorical Authority. My aim in this chapter is to examine the first of these, showing how an appeal to necessary moral truths can help us answer two important objections to it.

More precisely, in this chapter I am concerned with the following claim about moral properties:

Non-Reductivism. Moral properties are not reducible to purely descriptive properties.

I clarify and motivate this claim in §3.1. The rest of the chapter is devoted to answering the two most influential ways of objecting to it. Both of these objections rely on the fact that the moral ‘supervenies’ on the descriptive. I explain what this supervenience claim involves in §3.2, before outlining the first supervenience objection against Non-Reductivism in §3.3. I call this the ‘*Reductive* Supervenience Objection,’ and I respond to it in §3.4. I outline the second supervenience objection against Non-Reductivism in §3.5. I call this the ‘*Explanatory* Supervenience Objection,’ and I respond to it in §3.6 and §3.7. My response to both supervenience objections will appeal to necessary moral truths, so I briefly discuss the existence of such truths in §3.8. I conclude, in §3.9, that an appeal to necessary moral truths enables Robust Realists to retain the commitment to an irreducible moral reality.

3.1 Non-Reductivism and Descriptivism

What does commitment to Non-Reductivism involve, and with what claim does it most directly contrast? In this section I answer these questions, and in doing so I provide evidence for Non-Reductivism. The evidence is defeasible, for it relies on how things appear. It is for Non-Reductivism’s critics to object that things are *not* as they appear. If the objections can be rebutted, we are provisionally entitled to say that moral properties don’t reduce.

First, how are we to distinguish moral and descriptive properties? I said in §1.1 that descriptive properties are picked out by the descriptive predicates, like ‘... is human’ and ‘... is desired,’ whilst moral properties are picked out by moral predicates, like ‘... is right’ and ‘... is good.’ This approach is not uncontroversial, but it has the attraction of ensuring that the debate begins on an even footing.¹ After all, to draw the distinction in metaphysical terms would beg the question against the Non-Reductivist’s critics, who deny that the moral and the descriptive are metaphysically distinct. Maybe there are other ways of drawing the distinction that are equally dialectically friendly but, by approaching the properties via the predicates that pick them out, we allow that these predicates may or may not refer to just one type of property. The debate can then be had on even terms.

If moral and descriptive predicates refer to one type of property, then it will most likely be due to the truth of the following:

Descriptivism. Moral properties are reducible to purely descriptive properties.

Descriptivism is the most direct contrast with Non-Reductivism. You may ask why I treat descriptive properties as the potential reductive base for moral properties, rather than the (perhaps) more traditional natural or non-moral properties. With the natural, there are two reasons. One is that I suspect that there is room for a non-reductive form of moral naturalism. I won’t explore this view now – I discuss it in Chapter 5 – but we shouldn’t view moral naturalists as necessarily being hostile to the irreducibility of moral properties.² A second reason is that, in accepting Non-Reductivism, the Robust Realist is partly seeking to exclude theistic views that reduce, say, *being good*, to a supernatural property like *being loved by the gods*. Non-Reductivism is thus not exclusively against naturalistic reductions, for it is also against supernaturalistic reductions. The term ‘descriptive’ is therefore useful, for it broad enough to incorporate the supernatural as well as the natural. So, in short, the Non-Reductivist should deny both that moral properties are descriptively natural and that they are descriptively supernatural.

¹ It is used by critics of Non-Reductivism, like Jackson (1998) and Streumer (2008), so by operating with it we’re clearly not stacking the deck in favour of Non-Reductivism.

² The view I have in mind is mostly associated with Boyd (1988), Sturgeon (1988; 2006), and Brink (1984; 1989; 2001). I consider this type of view in Chapter 5, in which I also discuss the relation between the natural and the descriptive in more depth.

With the non-moral, it is just wrong to say that moral properties might reduce to non-moral properties. In this setting, attempts to reduce moral properties are meant to be ‘conservative,’ not ‘eliminative.’ That is, reductive metaethical claims do not aim to usher moral properties from the scene. They instead aim to capture moral properties via more basic properties in which they putatively consist. And something that counts as moral cannot consist in anything non-moral. This point is made by Mark Schroeder, who says that a “property cannot be both moral *and* nonmoral. If it is non-moral, then it is *not moral* after all” (2005: 9). By talking of the relation between the moral and the descriptive, rather than the moral and the non-moral, ensures that we avoid this unhelpful view of the reductive claim.

So, it’s best to talk of the descriptive instead of the natural or the non-moral. But you might worry that moral properties are just obviously descriptive. After all, ‘killing is wrong’ seems to describe killing as having the property of wrongness, just as ‘the table is brown’ describes the table as having the property of brownness. This is no longer an issue, however, if we clarify the precise sense in which Non-Reductivists deny that moral properties are descriptive. For the descriptive, as it is understood here, is firmly on the ‘is’ side of the ‘is/ought’ gulf. Whilst no one could deny that morality includes ought-y stuff, the most important question is whether this ought-y or directive part of morality can be had in purely descriptive terms.

Non-Reductivists say that it cannot, for it is far from clear that to direct is simply to describe in a certain way. Maybe there is a way to have direction in purely descriptive terms, but this is where the disagreement between Non-Reductivists and Descriptivists is located. Put another way, there is a similarity between ‘killing is wrong’ and ‘the table is brown’ in that both claims are truth-apt.³ But ‘killing is wrong’ is not *just* truth-apt. It is also ‘satisfaction-apt,’ as I will put it. It involves or entails a directive claim about how one ought to behave, and a moral agent can thus satisfy or fail to satisfy it. A reductive claim must capture the satisfaction-aptness of moral property ascriptions. Descriptivists and Non-Reductivists divide over the prospects of capturing satisfaction-aptness in purely descriptive terms.

³ I thus use ‘descriptive’ differently to Hare (1952), who denies that moral sentences are truth-apt. I set aside Hare’s prescriptivist view. The relevant debate is between those, like Jackson (1998) and Brown (2011), who think that the whole of morality can be captured on the ‘is’ side of the ‘is/ought’ divide, and those who think that it cannot.

This satisfaction-aptness certainly doesn't seem to be present in some potential reductions. If one attempted to reduce *being a wrong act* to *being an act that fails to maximise happiness*, then one would appear to lose satisfaction-aptness. After all, 'killing is wrong' appears to involve or imply direction that agents can satisfy or fail to satisfy. But this isn't clear of 'killing fails to maximise happiness.' This claim is truth-apt, but it does not appear to be satisfaction-apt. Other cases are more delicate, however. For instance, suppose that *being a wrong act* is reducible to something like *being an act that is against the commands of the gods*. Then the claim 'killing is wrong' equates to 'killing is against the commands of the gods.' This may look satisfaction-apt, for agents can satisfy or fail to satisfy a command. But careful reading suggests that it not, in fact, satisfaction-apt. For it does not *issue* a command, it only *describes* an act as having been commanded. To see the difference, compare a command like 'do not kill' with the description that 'killing is against the commands of the gods.'⁴ Clearly, the former is satisfaction-apt. But it is not clear that the latter is satisfaction-apt. I know how to obey a command that has been issued. I don't know how to obey a description of the fact that a command has been issued.⁵

So, there seems to be a key difference between the moral and the descriptive. Directive moral claims are satisfaction-apt, and it appears as though satisfaction-aptness cannot be had in purely descriptive terms. To clarify how this helps Non-Reductivism, I should make explicit the notion of reduction that I have in mind. In claiming that moral properties reduce to purely descriptive properties, I take it that most Descriptivists mean that moral properties are *identical* to some purely descriptive properties. They think that the set of moral properties is included in the set of purely descriptive properties.⁶ In contrast, Non-Reductivism (as I understand it) claims that moral properties are *not* identical to any purely descriptive properties.

⁴ If we see 'killing as wrong' as meaning 'do not kill,' we are dealing with an Expressivist rather than a Descriptivist view. That's a different kettle of fish, discussed in Chapter 2.

⁵ The point about commands applies to desires too. There's a difference between having a desire and a description of the fact that a desire is had. I know how to satisfy the former, but I don't know how to satisfy the latter.

⁶ This is the approach taken by Jackson (1998) and, if I read him correctly, Brown (2011). Schroeder (2005; 2007) has a different approach, but I set this aside as I think that the identity claim is what most Robust Realists are chiefly concerned to deny.

Having understood reduction as identity, we can be clearer about the way in which morality's satisfaction-aptness helps to make a case for Non-Reductivism. Consider the following part of Leibniz's Law:

Indiscernibility of Identicals. For any *a* and *b*, if *a* is identical to *b* then, for all properties *F*, *a* has *F* if and only if *b* has *F*.

From this highly plausible principle we can draw a test that enables us to assess whether two entities are metaphysically distinct. The test is just to see whether they have any different properties. If they do, the entities are indeed non-identical and therefore distinct. In the metaethical case, our question is thus as follows: is there any difference in the (second-order) properties of (first-order) moral and descriptive properties? According to Non-Reductivists, moral properties appear to have the (second-order) property of *being directive*, or *being satisfaction-apt*. Purely descriptive properties seem to lack this (second-order) property.⁷ If the appearances reflect reality, moral and descriptive properties will indeed have different properties. Given the test provided by Indiscernibility of Identicals, moral properties and descriptive properties will therefore be non-identical and distinct.

There is thus evidence for Non-Reductivism. But this evidence is defeasible, for it comes from *apparent* differences between the moral and the descriptive.⁸ It *appears* that morality involves satisfaction-aptness, and it also *appears* that we can't get satisfaction-aptness from the purely descriptive. But there may be arguments that make us think twice about the appearances. Argument may show that satisfaction-aptness is purely descriptive. Or it may make us revise our view of morality, so that we see the appearance of satisfaction-aptness as an illusion. In what follows, I look at the best ways of objecting to Non-Reductivism. If, as I'll argue, they fail to bite, then the evidence for Non-Reductivism will stand, and we will be (provisionally) entitled to maintain our belief in an irreducible moral reality.

⁷ I argue in Chapter 5 that the specifically *categorical* nature of moral direction is evidence for Non-Naturalism, but this is a different point to the one made here.

⁸ Enoch (2011: 82) judges that normative properties are irreducible by appealing to what he calls the 'just-too-different' intuition. This may be similar to the point made here, but Enoch does not specify a view of reduction. Seeing reduction as identity clarifies that the support for Non-Reductivism is not *bare* intuition. Instead the idea is that, in concord with Indiscernibility of Identicals, apparent differences between the moral and the descriptive provide evidence for the distinctiveness of those properties.

3.2 Moral Supervenience

The most influential objections to Non-Reductivism draw on the widely accepted fact that moral properties supervene on descriptive properties. The notion of supervenience is sometimes less than clear in discussions of these issues, so before I consider how ‘supervenience objections’ are meant to be a threat to Non-Reductivism, I’ll discuss moral supervenience itself.

The supervenience relation has been understood in various ways over the years. In its early uses, it was treated by some as a certain sort of metaphysical relation. Specifically, it was often seen as an asymmetrical dependence relation.⁹ But this is not how it is understood now. Instead, supervenience is standardly understood as a *modal* relation that holds between classes of properties. In other words, to say that a certain set of properties supervenes on some other set of properties is to say that those classes of property co-vary across worlds. Understood in these terms, the supervenience relation is neither symmetric nor asymmetric. Instead it is *non-symmetric*. That is, supervenience claims in themselves presuppose nothing about the character of any metaphysical link between the sets of properties with which it is concerned. Two sets of properties related by supervenience might be related symmetrically (for instance, by identity) or asymmetrically (for instance, by dependence), but this isn’t decided by the fact that one of them supervenes on the other. In sum, supervenience is a non-symmetric modal relationship that can hold between property classes.

Philosophers also draw a distinction between different *types* of supervenience, and this is partly because the worlds across which two or more sets of properties co-vary can be more or less wide-ranging. It might be that the relevant sets of properties are correlated within a limited set of worlds, for example, or it might be that the relevant sets of properties are correlated across all worlds. Fortunately, in the present context there is no need to flesh out every type of supervenience.¹⁰ Given that our only concern at the moment is the connection between moral properties and descriptive properties, just one type of supervenience relation is relevant to us.

⁹ McPherson (2012: 216) makes the same point.

¹⁰ Kim (1993) has a seminal discussion, and McLaughlin and Bennett (2014) give a broad survey of both classic and more recent literature.

Specifically, debate between Non-Reductivists and Descriptivists involves the following 'global' type of supervenience:

Global Supervenience. The *A*-properties globally supervene on the *B*-properties if and only if any two worlds that are exactly alike in their *A*-properties are also exactly alike in their *B*-properties.

Among metaethicists, it is widely (though not universally) accepted that moral and descriptive properties are related in this way.¹¹ I will thus take the following as common ground:

Moral Supervenience. Any two worlds that are exactly alike in their descriptive properties are exactly alike in their moral properties.

Moral Supervenience states that the moral properties and the descriptive properties co-vary across *all* worlds. Although it is not often given explicit argumentative support, its plausibility can be illustrated by an appeal to cases.

For instance, if worlds *w* and *w** are exactly alike descriptively, then the claim that Arthur acted wrongly in slapping Barry in *w* but that Arthur* acted rightly in slapping Barry* in *w** will seem highly peculiar. Such a claim appears to lack any motivation. It seems to be arbitrary and inconsistent, maybe *ad hoc*. Cases violating Moral Supervenience are thus rejected as impossible. Basically, then, it is counterintuitive to claim that there could be a moral difference between two worlds without there also being some relevant descriptive difference to explain it. Of course, if there were some relevant descriptive difference between *w* and *w**, then the existence of corresponding moral differences between these two worlds would be fine. Suppose that Arthur* performed his slapping in order to prevent Barry* from fainting, and that this was not the case with Arthur. There can of course be a moral difference between *w* and *w** in this sort of situation. But this doesn't entail a violation of the Moral Supervenience principle, for the two worlds are no longer descriptively alike. Without a relevant descriptive difference, there cannot be any corresponding moral difference.

¹¹ Sturgeon (2009) raises interesting doubts, but I think that there are strong responses to them – see Ridge (2007) and McPherson (2012). It might be that Moral Supervenience, as I've stated it, should be slightly reformulated in light of Sturgeon's discussion, but the relevant reformulations will make no difference to the arguments of this chapter. To keep things simple and clear, I therefore set them aside.

Moral Supervenience is not optional. Non-Reductivists are thus in trouble if, as some suggest, it makes problems for them. One can't escape any problem it creates by denying it.¹² One must make instead some other move.

3.3 Reductive Supervenience Objections

One way of challenging Non-Reductivism says that Moral Supervenience, when combined with certain other claims, *entails* Descriptivism. That is, some say that the modal link between the moral and the descriptive helps to establish the reducibility of moral properties to descriptive properties. Call this the 'Reductive Supervenience Objection.' In this section I outline the best way of framing this objection to Non-Reductivism. I respond to it in §3.4.

The most influential Reductive Supervenience Objection is due to Frank Jackson, but I will suggest that it isn't the most worrying. Jackson's version of the objection relies on a contentious metaphysical claim, and as a result has little or no dialectical bite against Non-Reductivism. Before coming to a more troubling version of the objection, however, we should consider Jackson's influential argument. He begins as follows:

Let *E* be a sentence about ethical nature ... Now each world at which *E* is true will have some descriptive nature: ethical nature without descriptive nature is impossible ... And, for each such world, there will be a sentence containing only descriptive terms that gives that nature in full. Now let *w*₁, *w*₂, etc. be the worlds where *E* is true, and let *D*₁, *D*₂, etc. be purely descriptive sentences true at *w*₁, *w*₂, etc., respectively, which give the full descriptive nature of *w*₁, *w*₂, etc. Then the disjunction of *D*₁, *D*₂, etc., will also be a purely descriptive sentence, call it *D*. But then *E* entails and is entailed by *D*. For every world where *E* is true is a world where one or other of the *D*_{*i*} is true, so *E* entails *D*. Moreover, every world where one or other of the *D*_{*i*} are true is a world where *E* is true, as otherwise we would have a violation of [Moral Supervenience]. Therefore, *D* entails *E* (1998: 122-123)

The upshot of these comments is that moral predicates and descriptive predicates are necessarily coextensive. But we need more to get us to Descriptivism, which is a claim about moral properties rather than moral predicates.

¹² Though see Harrison (2013).

Jackson offers the following principle in order to get us all the way to Descriptivism:

Necessary Coextension. If two predicates are necessarily coextensive, then they refer to one and the same property.¹³

Given that moral and descriptive predicates are necessarily co-extensive, and assuming that necessarily coextensive predicates refer to one and the same property, it follows that moral and descriptive predicates refer to one and the same property. Moral and descriptive properties are one and the same. Given that there is no serious way for descriptive properties to reduce to moral properties, the reduction obviously goes the other way. That is, the moral reduces to the descriptive. And this is Descriptivism. In sum, Jackson seeks to move from Moral Supervenience to Descriptivism, via Necessary Coextension. If he succeeds, then the claim that there exist irreducible moral properties must be rejected in favour of one or another reduction of the moral to the descriptive.¹⁴ So, how can Non-Reductivists reply?

A standard reply is to reject Necessary Coextension. Critics offer cases like the following:

Triangle. The predicates ‘... is a closed figure with three sides’ and ‘... is a closed figure with three angles’ are necessarily coextensive, but these predicates refer to different properties.¹⁵

Nine. The predicates ‘... is the cube root of 729’ and ‘... is the square of 3’ are necessarily coextensive, but these predicates refer to different properties.¹⁶

The properties picked out by these necessarily coextensive predicates are said to be non-identical, and thus distinct. If this is correct, it undermines Jackson’s argument by showing that moral and descriptive properties can differ despite the necessary coextension of those predicates that refer to them.

¹³ Jackson (1998: 125-128).

¹⁴ This assumes that we aim to retain moral truths and properties. We could abandon that and be Error Theorists or Expressivists, but I set this aside. Cf. Streumer (2008; 2011).

¹⁵ Shafer-Landau (2003: 91), Majors (2005: 488), FitzPatrick (2008: 199), and Kramer (2009: 209-210).

¹⁶ Oddie (2005: 149) and Parfit (2011 v.2: 297). Note that, whilst Oddie is committed to Non-Reductivism, he discusses this case in order to argue that it is unhelpful.

But how do defenders of Non-Reductivism argue that, in Triangle and Nine, the relevant predicates refer to distinct properties? The answer starts with the Indiscernibility of Identicals test mentioned in §3.1. Recall that, on this test, we can see whether two entities are identical by asking if they have the same properties. If they don't, they are non-identical and therefore distinct. The idea in the Triangle and Nine cases is just that the (second-order) properties of the (first-order) properties in question are not entirely alike, and that as a result those (first-order) properties are distinct.

In the case of Nine, for example, the property of *being the cube root of 729* and the property of *being the square of 3* each have different (second-order) properties. To illustrate, consider that one can investigate the cube root of 729 without investigating the square of 3. This means that the property of *being the cube root of 729* has the (second-order) property of *being investigable in isolation from investigation into the square of 3*. This is not true of the property of *being the square of 3*. Given the Indiscernibility of Identicals principle, it therefore seems as though *being the cube root of 729* and *being the square of 3* are distinct, for these (first-order) properties have distinct (second-order) properties. Equally fine-grained differences can be found for the Triangle case. This is why most Non-Reductivists reject Necessary Coextension. With this principle rejected, the door to understanding moral and descriptive properties as distinct – despite the necessary coextension of their predicates – remains open.

If this is right, then Non-Reductivism still stands. However, there is a worry about this response. Some deny that Triangle and Nine count as counterexamples to Coextension, for they don't share the intuition that the predicates in these cases refer to distinct properties. Instead, they are inclined to judge that '*... is a closed figure with three sides*' and '*... is a closed figure with three angles*' refer to one property – *being a triangle*. Similarly, they judge that '*... is the cube root of 729*' and '*... is the square of 3*' refer to one property – *being the number 9*. The core charge is that Triangle and Nine conflate properties with the contents of predicates, and that they thus fail to support Non-Reductivism.¹⁷ Of course, not everyone shares the intuitions required to underpin that charge. Non-Reductivists typically don't, which is why they appeal to Triangle and Nine in the first place.

¹⁷ Jackson (1998: 125-128) and Streumer (2008; 2011).

This is the cause of my concern about Jackson's way of developing the Reductive Supervenience Objection – it seems that, when the debate about supervenience and reduction is framed in this way, it results in a stalemate.¹⁸ Descriptivists and Non-Reductivists both end up defending arguments that fail to be dialectically effective, for they rely on intuitions that their opponents – the very people they are trying to convince – lack. Specifically, intuitions concerning highly abstract issues pertaining to the Necessary Coextension principle. Given that such differences in intuition tend to be very difficult to adjudicate, it is likely that Jackson's argument will ultimately offer little assistance as a way of settling the debate about supervenience and reduction in metaethics. True, it may be that there is a way out of the impasse. But it may also be that we have to move on from Jackson's supervenience objection to move forward. I suspect that this is the case, so I won't discuss Jackson's objection further. However, there is a way of framing the Reductive Supervenience Objection that presents a more worrying threat to the claim that there exists an irreducible moral reality.

This threat comes from an argument made by Campbell Brown, an argument that goes as follows:

- (1) Moral properties supervene on descriptive properties.
- (2) If moral properties supervene on descriptive properties, then all non-descriptive moral properties are redundant in the sense that they do no work in distinguishing possibilities.
- (3) No properties are redundant.
- (4) So, all moral properties are descriptive properties.¹⁹

Before considering how Brown's version of the Reductive Supervenience Objection improves upon Jackson's, let's consider its premises in more detail. As we'll see, it only makes a few claims, each of which seems highly plausible.

¹⁸ This is also suggested by Suikkanen (2010: §3). Suikkanen seeks a way to move forward in the debate about Jackson's objection. This is a legitimate avenue to explore, though I am not persuaded that it is ultimately a fruitful one. Given limitations of space, I prefer to focus on Brown's more recent version of the Reductive Supervenience Objection. I discuss his argument below.

¹⁹ Brown (2011: 210). He goes on to modify the argument in light of a potential objection, but the modification makes no difference to my discussion so I set it aside.

The claim made by (1) is that moral and descriptive properties have a certain modal relationship, and Brown is clear in his discussion that he has Global Supervenience in mind. In other words, the first part of this Reductive Supervenience Objection is simply Moral Supervenience. I have already considered this principle in §3.2. It is widely accepted and I take it to be common ground, so I won't discuss it further here. The claim made by (2) is that Moral Supervenience entails the redundancy of moral properties, where redundancy is interpreted as follows:

Redundancy. A set of properties is redundant just in case it makes no contribution to distinguishing possibilities.

Brown thinks that moral properties fail this test. That is, the set of moral properties makes no contribution to distinguishing possibilities. To see this consider that, if *A*-properties supervene on *B*-properties, two worlds that are alike in their *A*-properties will thus be alike in their *B*-properties. In other words, the worlds will be indistinguishable in respect of these properties. So, if redundancy consists in failure to make a contribution to distinguishing possibilities, all supervening properties will be redundant. Moral properties supervene, and are therefore redundant in the relevant sense.

The claim made by (3) is that there are no redundant properties. Brown's rationale for this position is a "maxim of ontological parsimony: posit only so many properties as are required to distinguish possibilities" (2011: 212). This is just a specific version of 'Ockham's Razor' – a widely accepted norm that tells us to accept the simpler of two equally successful explanations. Now, if moral properties are redundant, and if there are no redundant properties, then we are only entitled to posit moral properties if we reduce them to the descriptive properties on which they supervene. We must therefore see moral properties as purely descriptive properties. This is (4), which is just Descriptivism. In short, Brown seeks to move all the way from Moral Supervenience to Descriptivism, via his Redundancy principle.

The attraction of framing the Reductive Supervenience Objection in this way is that it avoids the risk of stalemate that arises from Jackson's reliance on Necessary Coextension, a principle about which people differ in intuition. Brown's argument relies on no such controversial principle to bridge the predicate-property gap. It is a directly metaphysical case for Descriptivism, and relies on just three principles that are apparently very plausible: Moral Supervenience, Redundancy, and a version of Ockham's Razor.

This way of formulating the Reductive Supervenience Objection is therefore the deeper threat to irreducible moral properties. Defenders of Non-Reductivism need a way to avoid or overcome this threat. I offer them a way of doing this in the next section, in which I argue that one of Brown's principles is not as plausible as it seems. It needs to be modified to be defensible, and once modified it opens the door to a Non-Reductivist view.

3.4 Necessity and Redundancy

One way of responding to Brown's supervenience objection would be to show that irreducible moral properties are not redundant, that they do contribute to our being able to distinguish possibilities.²⁰ Another way of responding to the objection would be to accept that moral properties are redundant in Brown's sense, whilst also arguing that there are redundant properties in exactly this sense.²¹ These are both interesting ways to go, but I intend to offer a different sort of response to Brown. My preferred way of responding is to cast doubt on redundancy as Brown understands it.

Recall that Brown's supervenience objection to Non-Reductivism relies on the following conception of redundancy:

Redundancy. A set of properties is redundant just in case it makes no contribution to distinguishing possibilities.

Now, this principle has to be modified in some way, for it cannot be true as it stands. To see this, consider mathematics. The pure and basic truths concerning mathematical entities hold necessarily. For instance, it is true at all possible worlds that the number nineteen has the property of being prime. In general, all worlds are indistinguishable in respect of such pure and basic mathematical truths. But should we therefore conclude that the mathematical is redundant? This, it seems to me, would be an extremely strange basis for that claim. Even if there are reasons to be suspicious of abstract mathematical entities, the fact that truths about them do not help us to distinguish possible worlds is not one of them. I'm not saying that mathematical entities exist. I'm just saying that the supervenience of the mathematical would be an odd basis for treating mathematical entities as redundant.

²⁰ This is how Enoch (2011: 139-140) responds to Brown.

²¹ This is how Wielenberg (2011) responds to Brown. For a reply, see Brown (2012).

To clarify, what is odd is the idea that redundancy automatically falls out of mathematical supervenience. Maybe mathematical entities are dispensable, and thus redundant, but it will take a lot of work to expose this fact. Specifically, one must show that the best scientific theories can do without mathematics. Perhaps this can be done, but actually doing it will be extremely difficult. After all, even one of the best known and most detailed stabs at doing this sort of technical work – namely, Hartry Field’s effort to show that mathematical abstracta are dispensable in Newtonian mechanics – is explicitly incomplete.²² So, my claim is not that there are mathematical entities. I’m simply agreeing with Mark Colyvan’s view that mathematical nominalists have to take the ‘hard road,’ like Field, for there is unlikely to be any ‘easy road’ available.²³ And, given that it would be unacceptably strange to regard mathematical entities as dispensable and redundant on the basis that they don’t help us to distinguish worlds, we should reject the formulation of Redundancy on which Brown’s argument relies.

Again, this is not to claim that there are mathematical properties. I am neutral on that score. The point is just that it is a mistake to regard this debate in the philosophy of mathematics as being settled just by the contribution that the truths concerning mathematical properties make to our distinguishing between possible worlds. It is clear that they make no contribution here, for pure and basic mathematical truths are the same at all worlds. Now, one might think that this just points us toward a fix for Brown’s argument. After all, it is pretty well known that the necessity of mathematical truths makes them a special case within the supervenience debate. They are a special case because truths concerning mathematical properties hold as a matter of necessity, and this ensures that *every* class of property supervenes upon them. Two worlds that are exactly alike in their geographical or psychological properties, for example, will also be exactly alike in their mathematical properties. And that’s simply because *all* possible worlds are exactly alike with respect to their mathematical properties.²⁴

²² Field (1980).

²³ Colyvan (2010) argues that three influential, attractive, and representative attempts to take the easy road in the end require the success of a hard road strategy. As he makes clear, the claim is not that there is *definitely* no easy road. It’s just that the prospects look dim, given that the best shots at providing one fail in the relevant way.

²⁴ Cf. Williamson (2001: 626).

Things are different with most other supervenience claims, for to say that the class of *A*-properties supervenes on the class of *B*-properties is typically to say that those classes of property necessarily co-vary across worlds. And there is no co-variance in the mathematical case, for the basic mathematical truths don't change – they are necessary. The mathematical case thus differs from most other cases of supervenience. So, even though Redundancy has strange results for the mathematical, the fact that this is a special case means we can probably modify it to avoid those results:

Modified Redundancy. If the truths pertaining to a set of properties are all contingent, then that set of properties is redundant just in case it makes no contribution to distinguishing possibilities.

Note that this modification of Redundancy is not *ad hoc*. It is a response to the widely accepted fact that the necessity of the mathematical makes it a special case in the supervenience debate. So, even though Redundancy is a flawed test for redundancy, Modified Redundancy has more going for it.

Perhaps, then, this Reductive Supervenience Objection still works against Non-Reductivists. That is, Brown's argument might succeed when Redundancy is replaced by Modified Redundancy. However, the matter is more complex than Descriptivists might hope, for Modified Redundancy opens the door to a Non-Reductivist view that accepts necessary moral truths. Basically, if we say that some moral truths, like the pure and basic mathematical truths, are fixed across all worlds, then Non-Reductivism is left untouched by a version of Brown's argument that has undergone the modifications required to avoid odd results in the mathematical case. So, accepting necessary moral truths provides Non-Reductivists with a way forward. The idea that there are necessary moral truths is controversial, of course.²⁵ Many moral truths are quite clearly contingent. *It was morally impermissible for me to punch Perkins in the face.* This is true, but not in all worlds. There are worlds in which we don't exist, worlds in which we do exist but I didn't punch Perkins in the face, worlds in which the punch was the only way to avoid war, and so on. So, even if it was impermissible for me to punch Perkins in the face, this is evidently a contingent moral truth.

²⁵ Various metaethicists – like Parfit (2006), Enoch (2011), and Scanlon (2014) – defend it, but it is still up for debate.

Perhaps, then, it will be unwise for the Non-Reductivist to rely on necessary moral truths in answering Brown's supervenience objection. To see whether this is the case, we need to hear more about the idea that there are necessary moral truths. I will delay that discussion, however, so that I can examine it more fully later on. I argue in §3.8 that we can at least reasonably believe in necessary moral truths, and in Chapter 4 I provide a more detailed model of such truths. On my view, there are moral *norms* whose modal status is not limited by any contingent presuppositions or facts, just as there are basic mathematical truths whose modal status is not limited by any contingent presuppositions or facts. I will get to this in due course, but first let's consider the second supervenience objection. As I'll explain, the answer to this objection requires necessary moral truths as well.

3.5 Explanatory Supervenience Objections

Reductive Supervenience Objections seek to establish Descriptivism. The second supervenience objection that I'll consider is more modest, for it doesn't try to establish a specific theory. It just says that Non-Reductivists cannot explain the modal link between the moral and the descriptive. In short, if moral properties are irreducible, then moral supervenience is a mystery.

Call this the 'Explanatory Supervenience Objection.' In some early discussions of this objection, it can be hard to tell exactly how it is meant to work. Consider the following, from J.L. Mackie:

What is the connection between the natural fact that an action is a piece of deliberate cruelty ... and the moral fact that it is wrong? ... The wrongness must somehow be 'consequential' or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this 'because'? (1977: 41).

Mackie finds Moral Supervenience mysterious, or 'queer.' But this doesn't amount to an objection to Non-Reductivism. It seems more like a call for information. Mackie says nothing to show that Non-Reductivists cannot just answer this call, and nothing to expose any special mystery for that view.²⁶

²⁶ Mackie does go on to explain supervenience via our subjective responses to the natural world. However, he does not develop a potential account for Non-Reductivists, so he has nothing with which to compare his subjectivist explanation.

So, we must look elsewhere to find a version of the Explanatory Supervenience Objection that has bite against the irreducibility of moral properties. Simon Blackburn (1993b; 1993c) has a more rigorous version of it. His discussion is hard to interpret, but one reading is that he thinks it a conceptual truth that there can be no moral changes without relevant descriptive changes, and that Non-Reductivists are unable to explain this conceptual truth.²⁷ Whether or not this is what Blackburn had in mind, it is a mistake to suppose that it is a deep problem for Non-Reductivists. It plausibly is a conceptual truth that moral properties supervene on *some* descriptive properties – if someone were to deny this we’d regard them as conceptually deficient – but, as David Enoch (2011: 149) notes in this context, conceptual necessities do not cry out for explanation. So this way of understanding the Explanatory Supervenience Objection lacks serious bite.

However, there is an Explanatory Supervenience Objection that has bite. Whilst it is a conceptual truth that moral properties supervene on some descriptive properties, it is not a conceptual truth that moral properties supervene on the particular set of descriptive properties they in fact supervene on.²⁸ Suppose, for example, that the moral property of *being a wrong act* globally supervenes on the purely descriptive property of *being an act that fails to maximise happiness*. If this is the case, then it is the case in all possible worlds. And that means that many things that seem possible, and are defended by philosophers, are not in fact possible at all.

For instance, if *being a wrong act* supervenes on *being an act that fails to maximise happiness*, then (contra Kantianism) there is no world at which an act is wrong just because treats another merely as a means. The fact that it treats another merely as a means may be involved in its being wrong, but only if it somehow results in its failure to maximise happiness. In short, we must explain why the combination of *being a wrong act* and *being an act that treats someone merely as a means* is impossible rather than just non-actual. That is, we must answer the following question: why do moral properties supervene on the set of descriptive properties they in fact supervene on? All metaethical theories must answer this question, but some say that Non-Reductivism is in an unusually bad position on this score.

²⁷ Dreier (1992) helpfully interprets and assesses Blackburn’s version of the argument.

²⁸ Cf. Enoch (2011: 142) on the distinction between ‘general’ and ‘specific’ supervenience.

After all, the Descriptivist can explain this modal relation in terms of a metaphysical relation – identity – that they say holds between moral and descriptive properties. And the Expressivist can explain it by saying that the purpose of moral judgement is to guide desires and choices about descriptive parts of reality in a fully consistent way.²⁹ But what explains the modal relation if moral and descriptive properties are metaphysically distinct? This is a problem that seems unique to Non-Reductivism, and it is the core of the Explanatory Supervenience Objection as I understand it here.

Note as well that, according to the defenders of the Explanatory Supervenience Objection, we cannot reply by appealing to some brutally necessary connection between the descriptive and the moral. Building on traditional Humean dogma, according to which there can be no necessary connections between distinct existences, critics of Non-Reductivism hold that one cannot say that it is just a basic fact about reality that moral and descriptive properties are modally linked in the way they are. At worst, they see it as impossible for there to be such brute necessary connections. At best, they maintain that the commitment to such connections counts significantly against a view.³⁰ I discuss this in §3.7, but first I'll outline an answer the Explanatory Supervenience Objection on Non-Reductivism's behalf.

3.6 Necessity and Norms

My preferred answer to the Explanatory Supervenience Objection draws on necessary moral truths or, specifically, norms. This is not a novel claim. Torbjörn Tännsjö holds that such norms can answer the supervenience question, but the idea is developed in depth by Enoch.³¹ In this section I outline how we can explain supervenience by appeal to necessary moral norms. In the next section I defend the explanation against key objections to it.

²⁹ Blackburn's (1993c: 137), though this explanation is problematised by Dreier (2015).

³⁰ Hume (*Understanding*: §VII) and McPherson (2012: 217). Wilson (2010; 2014; 2015) argues that this view is currently unmotivated – the motivations that Hume gave for it are implausible, and other potential motivations are also unconvincing. She thus suggests that we should not give this Humean idea such a big role in the dialectic. Hills (2009) says that realists might like to try dealing with their supervenience problem by appeal to constant conjunction and not necessary connection, but this seems to me a drastic move.

³¹ Tännsjö (2010: 47) and Enoch (2011: 143-146). Cf. Skarsaune (2015).

Before moving to the moral case, it is helpful to flesh out the core explanatory move by considering a non-moral case of supervenience that can be explained via the existence of various norms. Enoch appeals to the supervenience of *legal drinking status* on *age*. The idea is that, within a legal jurisdiction and other things being equal, if two people are the same age then either both of them are permitted to purchase alcohol or neither of them are permitted to purchase alcohol.³² This is obviously explicable via relevant legal norms: if the law is that one must be eighteen or over to purchase alcohol, then two eight-year-olds are both not permitted to do so and two eighty-year-olds are both permitted to do so. The legal norms mean that differences in legal drinking status co-vary with differences in age.

In other words, legal drinking status *supervenies* on age. As this is happening in a legal jurisdiction, the co-variance of legal drinking status on age is pretty localised. But, as Enoch points out, there is no mystery as to *why* legal drinking status supervenes on age. We can explain this just by appealing to the existence of certain legal norms that determine legal drinking status within a given legal jurisdiction. And this provides us with a general model for explaining cases of supervenience; the supervenience of *A*-properties on *B*-properties is explicable via norms that link one and the other. Call this 'The Norms Approach' to explaining supervenience. This model won't apply in all cases – that is, we will probably have to use models other than The Norms Approach to deal with certain other cases of supervenience. But all that matters here is its applicability to the moral case.

The idea in the moral case is fairly similar to that of the legal case. Just as there are legal norms that hold within a certain legal jurisdiction, there are moral norms hold within a certain moral jurisdiction. A moral norm directs a moral agent to behave in a certain way. Perhaps the major difference between the moral and legal case is the fact that, whilst a legal jurisdiction is modally restricted in certain ways, morality's jurisdiction is modally unrestricted. To use Enoch's helpful way of putting it, moral norms have a 'modally maximal jurisdiction.'³³ They hold at all possible worlds.

³² We need the 'other things being equal' qualifier because the law may involve certain subtleties. For instance, it may say that no one under 18 can buy alcohol, except 17 year olds who are with an adult and eating a meal. For clarity, we can set such subtleties aside.

³³ Enoch (2011: 145).

It is important to the success of the explanation of supervenience in the moral case that moral norms have a modally maximal jurisdiction. The supervenience in the legal case might be pretty localised but, as we've seen, Moral Supervenience is 'global.' Moral and descriptive properties co-vary across all worlds, so if this co-variance is explicable in terms of moral norms, the norms had better hold across all worlds. Enoch does not offer a specific way of modelling the modally maximal necessity that this approach requires of moral norms; he is neutral on this score. I consider this in the Chapter 4, in which I give an account of the necessity in Robust Realism.

But let's not get ahead of ourselves. First we must clarify how the explanation is actually meant to work. Enoch offers a nice illustration:

Suppose that some fairly basic version of utilitarianism is true. Then the relevant moral norm states, roughly, that an action is wrong if and only if there is an alternative action (or inaction) that could be opted for, and that had a larger (perhaps expected) utility value. Given that this is the content of the norm, no mystery remains about the supervenience of wrongness on the (perhaps expected) utility values of the relevant action and its alternatives (2011: 144).

So, if someone asks you why moral properties globally supervene on the specific set of descriptive properties that they in fact globally supervene on, you can just direct them to the relevant moral norms. It is those norms that link up a specific set of moral and descriptive properties and, because they hold necessarily, the link is the same across every world. Enoch's example uses a simple version of utilitarianism, but The Norms Approach could be adapted to any potential view of the content of necessary moral norms.

I said above that The Norms Approach is not the only explanatory model available to us. For instance, it is sometimes best to explain the supervenience of *A*-properties on *B*-properties by maintaining that there is an asymmetric dependence relation linking those classes of property. The modal relationship arises because of the fact that the *A*-properties are 'higher level' properties that metaphysically depend on 'lower level' *B*-properties. Call this 'The Levels Approach.' Some think that this model should be applied to the moral case. So, before considering an objection to The Norms Approach, it is worth clarifying why I prefer it to The Levels Approach.

After all, The Levels Approach is popular with Non-Reductivists. A number of versions of it have been proposed, often involving different dependence relations. One might go for ‘constitution’ or ‘realisation,’ for example, or ‘grounding,’ or ‘making,’ or some other relation.³⁴ In general, however, the claim is that there are lower level descriptive properties that in some sense ‘give rise to’ higher level moral properties, and that it is this metaphysical relation that explains the modal relation between the moral and descriptive. So, why can there be no moral difference without a descriptive difference? Because the moral properties are constituted by descriptive properties, or grounded in such properties, or whatever. They are modally linked because they are metaphysically linked. There is thus no moral difference without a descriptive difference, for the moral level somehow ‘emerges’ from the descriptive level. The Levels Approach thus seems to explain supervenience without commitment to necessary moral norms.

The devil is in the detail, and I don’t have room for details. But let me explain why The Norms Approach is more promising, or at least more basic, than The Levels Approach. Our rationale for Non-Reductivism is the apparent difference between direction and description – this is evidence for distinctiveness of moral and descriptive properties. My worry about The Levels Approach is that it is not clear how robustly irreducible moral direction can *emerge* from description, given that these seem to be very different things. That is, there appears to be a deep discontinuity between directive and descriptive properties. So, without something else to link them – a norm, say – it is hard to see how the directive *emerges* from the descriptive. Indeed, many suppose that there must be such links in other cases in which we might appeal to higher and lower level properties. For instance, we may need bridge laws to link the physical and the chemical, or the neurological and the psychological. So, it is plausible that we need bridge laws or norms within The Levels Approach, otherwise it will have nothing to say about the emergence of direction from description. Thus, even if there is an element of truth in The Levels Approach, it still requires appeal to norms. Personally, however, I find the image of levels unhelpful. I prefer to focus directly on necessary moral norms, leaving levels to one side.

³⁴ Shafer-Landau (2003) talks of *constitution*. Dancy (2004a; 2004b), Huemer (2005), and Wielenberg (2014) talk of *making*. Wedgwood (2007) talks of *realisation*. Audi (2013) talks of *grounding*, as does Väyrynen (2013), though he does not advocate it.

Put another way, given the deep discontinuity between direction and description, it is hard to see how appeal to some dependence relation alone can constitute a real advance for the Non-Reductivist. After all, the accusation against the Non-Reductivist is that they make a mystery out of morality's supervening on the descriptive. And, without bridging norms, The Levels Approach simply creates another mystery; that of direction's emergence from description. It is unclear how this emergence can happen without norms to bridge the gap between the moral and the descriptive levels. Thus, even if one appeals to levels, one must also appeal to norms. Moreover, the emergence question does not arise for moral norms that count as fundamental – fundamental entities don't emerge from anything. Thus, whatever problems it might face, The Norms Approach that I favour is more promising, or more basic, than The Levels Approach that many favour.

So, at present our best chance of satisfactorily responding to the Explanatory Supervenience Objection involves an appeal to necessary moral truths or, specifically, norms. The need to get clearer on the claim that there are such truths or norms is thus becoming more pressing. I will be coming to it very soon, I swear. But first I would like to examine a way of putting pressure on The Norms Approach as a genuine *explanation* of the modal connection between moral and descriptive properties. In doing so I will strengthen this way of answering the Explanatory Supervenience Objection.

3.7 Norms and Explanations

Suppose we agree that the property of moral wrongness supervenes on the descriptive property of failing to maximise utility. You might object to The Norms Approach by saying that, if someone requests an explanation for why this is true, they are unlikely to be satisfied by the claim that it is true because actions are wrong just in case they fail to maximise utility. At first glance, you might be inclined to think that this is no explanation at all. According to Erik J. Wielenberg, for example, this is nothing more than a restatement of what is supposed to have been explained.³⁵ In short, whilst we were seeking to *explain* the modal link between the moral and descriptive, all The Norms Approach appears to do is *insist* that there is such a modal link. It rephrases it via norm-talk, but to rephrase is not to explain.

³⁵ Wielenberg (2014: 22-23).

Another way to put the problem is to suggest that, in response to a call for explanation, the Non-Reductivist just says that there is a brutally necessary connection between distinct existences. The distinct existences in this case are moral properties like *being a wrong act* on the one hand, and descriptive properties like *being an act that fails to maximise utility* on the other. And critics of Non-Reductivism and The Norms Approach maintain that it is a problem if this is all that can be offered. They either deny that there can be brutally necessary connections between distinct existences, or they suggest that a commitment to such connections counts significantly against a theory. At best, then, it counts significantly against the commitment to Non-Reductivism if it ‘explains’ supervenience just by appealing to a brutally necessary connection. For that is no explanation at all.

But, properly interpreted, The Norms Approach *does* explain. For norm-talk doesn’t simply rephrase supervenience-talk. To see this, recall that supervenience is a modal relationship that holds between classes of property. A norm, however, is something else. I discuss norms in more detail in Chapter 4, but here we can say that norms are entities that direct agents in certain circumstances or kinds of circumstance. A legal norm is an entity that directs you not to drink if you are under 18. A moral norm is an entity that directs you not to kill just for fun. In short, norms are entities that govern the behaviour of an agent. It is by appealing to such entities that we are able explain cases of property co-variance. It would certainly be problematic for The Norms Approach if this very instance of co-variance were invoked to explain this very instance of co-variance. That would fail to explain, but that is not what’s happening. Instead, we appeal to a distinct entity – a norm – and use that to do the explanatory work.

Now, in certain cases, the norms can themselves be seen in more basic terms. For instance, there is no doubt a deeper story to tell about even the most basic legal norms, one that draws on our conventions and institutions to explain how such social entities come to be. It may be hard to find the true story, but we know that there is one. But the fact that there is a deeper story about legal norms does not mean that such norms do no explanatory work, for at a certain level of inquiry it is quite legitimate to draw on such norms in explaining the modal link between legal drinking status and age. There is a deeper story about the nature of legal norms, but such norms do good explanatory work at a non-fundamental level of inquiry.

In other cases, however, the norms cannot be understood in more basic terms. Robust Realists say that this is so in the moral case, they see the most basic moral norms as existing independently of conventions and institutions. (As I'll discuss in subsequent chapters, they see moral reality as mind-independent.) So, unlike the basic legal norms, there is no deeper story to tell for basic moral norms. But the idea is the same. We draw on the existence of certain distinct entities – fundamental moral norms – to explain the co-variance of moral and descriptive properties. Assuming that there is no deeper story to tell about moral norms, this is the level of inquiry at which the explanation of Moral Supervenience bottoms out.³⁶ So, by accepting fundamental moral norms as independent entities that link moral and descriptive properties, we explain supervenience in a way that is more than a restatement of the phenomenon that we were trying to explain. This is how The Norms Approach does genuine explanatory work.

You might worry that The Norms Approach is still unsatisfying because it still posits necessary connections between distinct existences. It is true that it involves a brute necessity, but this isn't a major problem.³⁷ The main legitimate worry about necessary connections between distinct existences – beyond dogmatic Humean prejudice – is that they fail to explain. And, as I have said, The Norms Approach does explanatory work. It involves a certain sort of bruteness, for it relies on fundamental moral norms. But this is not a bruteness that involves failure to explain. It is a bruteness that is just inevitable, given the way explanation works. Most explanations bottom out eventually. A non-fundamental moral norm may of course be explained via deeper moral norms, but ultimately we reach fundamental moral norms, which cannot be had in more basic terms. This makes them brute, but in an unobjectionable way. For the claim is not that supervenience itself is inexplicably brute. Instead the claim is that there exist fundamental ('brute') moral norms that can explain supervenience. And these fundamental norms are like fundamental laws of nature, in that to ask why they are as they are is to fail to understand what 'fundamental' means.

³⁶ You might deny that there are fundamental moral norms that cannot be understood in more basic terms. But that's not an *objection* to Non-Reductivism as I understand the view, it's just a bare-faced denial of it. We do need a debate about the existence of fundamental moral norms, but this will be a separate debate from the one we're having at present.

³⁷ Cf. Enoch (2011: 147).

In sum, it is a mistake to attack The Norms Approach on the basis that it posits necessary connections between distinct existences. Whilst it involves brute necessity, it is not a problematic aspect of the view. I thus claim that The Norms Approach stands as a response to the Explanatory Supervenience Objection against Non-Reductivism. Now, this means that I have drawn on the existence of necessary moral truths or norms in my responses to both of the supervenience objections I have considered. To ensure that these responses really are workable ways of rebutting efforts to undermine Non-Reductivism, we thus need to show that it is at least reasonable to believe that some moral truths or norms hold as a matter of necessity. In the next section, I will argue that this can be reasonably believed.

3.8 Necessary Moral Truths

Are there any moral truths that hold at all possible worlds? Maybe. That's as strong an answer as I'm willing to give in this section. I won't say that there definitely are necessary moral truths, I'll simply argue that it is not unreasonable to believe in them. This may seem rather modest, but I will explain how it is enough to ensure that we are provisionally entitled to retain our belief in the irreducibility of moral properties. After all, if we can show that it is at least reasonable to believe that there are necessary moral truths, then we will have a tenable reading of Non-Reductivism that survives the Reductive and the Explanatory Supervenience Objections. This will mean that argument has (so far) failed to overturn or outweigh the defeasible evidence in favour of the existence of an irreducible moral reality.

Now, some moral propositions may be thought to hold as a matter of necessity just because they are conceptual truths that help to structure moral thought and discourse itself. Take, for example, the idea that moral judgement is universalisable and the idea that 'ought' implies 'can.' These are at least plausible candidates for conceptual necessity. Whilst there are debates and controversies about them, even philosophers who question these claims will agree that there are things to be said in favour of seeing them as conceptual truths. I take it, however, that Non-Reductivists need something more substantive than universalisability or 'ought' implies 'can.'³⁸

³⁸ Cuneo and Shafer-Landau (2014) argue there are some substantive moral truths are conceptual truths. I address their view in Chapter 4.

After all, even if these claims count as moral truths, they are not very interesting moral truths. And, even if they are in some sense norms, they are not the sort of norms that help us to explain supervenience. So we need to look for substantive moral norms that can be reasonably seen as necessary. And, as I said in §3.4, this might be tricky. After all, many of the most everyday moral truths or norms are contingent. Here is a moral truth: *it was morally impermissible for me to punch Perkins in the face*. This may be true, but it is not true in all possible worlds. There is a world, for example, in which it was morally obligatory for me to punch Perkins in the face, for it was the only way to prevent a nuclear war. Even if we ask for a more general norm in virtue of which it was impermissible for me to punch Perkins in the face, we'll potentially get something contingent: *it is (pro tanto) impermissible to punch people in the face*. This is true in our world, but not in a world whose inhabitants like to be punched in the face.

However, the Non-Reductivist does not need every moral truth to be necessary. For my responses to both of the supervenience objections to work, we just need one or more necessary moral norms. So the fact that certain moral norms are contingent is fine, for what matters is that there is at least one respectable candidate for moral necessity. And there is. Ask yourself *why* my punching Perkins is impermissible. A plausible answer is that, in doing so, I deliberately caused undeserved and undesired pain that wasn't required as a means to any valuable end. For brevity, let's use the term 'T-Pain' to refer to undeserved and undesired pain that is not required as a means to any valuable end.³⁹ Here is a substantive moral norm: *it is morally impermissible to deliberately cause T-Pain*. This claim, I submit, is a credible candidate for a moral norm that holds at all possible worlds.

To assess whether it really is a necessary moral norm, we can try to imagine worlds at which it is obligatory or permissible to deliberately cause T-Pain. If we can assume that there is at least a rough correlation between conceivability and possibility, then the ability to imagine such a world would indicate that even this highly credible candidate for moral necessity is merely contingent. If we cannot imagine such a world, then this is (defeasible) evidence for there being at least one necessary moral truth.

³⁹ The 'T' is for 'Transworld' – my claim will be that Non-Reductivists may reasonably take it as a necessary truth that it is impermissible to deliberately cause such pain.

I suspect that different people will feel differently about this. But if at least some sincere and reasonable judges struggle to imagine a world at which deliberately causing T-Pain is obligatory or permissible, then the view that at least one moral norm is necessary can be reasonably held. I personally struggle to imagine a world at which it is either obligatory or permissible to deliberately cause T-Pain, and I think that there are others who are pretty much as sincere and reasonable as me who will struggle too.

I can imagine worlds at which causing *pain* is either obligatory or permissible. Indeed, ours is no doubt such a world. It may be that justice sometimes requires that we cause pain, for example. Punishment is rarely free of some sort of pain, after all. And it is often the case that causing pain is required to bring about a valuable end, such as when you put yourself through painful physical training to compete at a sporting event. So ours is a world in which causing pain is sometimes obligatory or permissible. I can even imagine worlds in which causing pain is *always* obligatory or permissible, for I can imagine worlds in which pain is only ever caused deliberately when it is deserved or desired or required as a means to a valuable end. However, none of these are examples of causing T-Pain, for to cause T-Pain is to cause undeserved and undesired pain that is not required as a means to any valuable end. It is the moral status of causing T-Pain that matters here, so we need to ask what its moral status is across worlds.

I can imagine worlds in which people *think* that causing T-Pain is obligatory or permissible. In some such worlds, people are misinformed about non-moral facts relating to the nature of pain. In others, they know these facts but have a warped sensibility. But the possibility of worlds in which people *think* that causing T-Pain is obligatory or permissible does not entail that it is actually so in such worlds. Some take morality to be mind-dependent in this way, but this has not been shown at this point in the dialectic and thus cannot be relied upon in rejecting necessary moral truths. I can also imagine possible worlds in which the impermissibility of causing T-Pain is in an interesting sense *irrelevant*. For instance, there are no doubt worlds whose inhabitants are incapable of causing pain, for example, or incapable of feeling it. The inhabitants of these apparently possible worlds would no doubt find it very difficult to understand or be concerned about the claim that it is impermissible to deliberately cause T-Pain.

However, the idea that causing T-Pain is morally impermissible is not thereby false in those worlds. It's simply irrelevant to the lives of its inhabitants. This is fine, and there may be moral norms – even necessary ones – that are irrelevant in our world. The *relevance* of a moral norm may be contingent, but this does not mean that such norms are *themselves* contingent. Thus, someone who is sceptical about the existence of moral norms that hold at every world must do more than just point to worlds at which the norms that seem most plausible for our world would fail to be relevant.

I'm struggling to conceive of a world at which causing T-Pain is obligatory or permissible. To the extent that conceivability determines or is at least correlated with possibility, I take this as (defeasible) evidence for the claim that there are no possible worlds in which it is obligatory or permissible to deliberately cause T-Pain.⁴⁰ On that basis, I am prepared to tentatively suggest that it is morally impermissible to cause T-Pain in all possible worlds. And, if this is right, there is at least one claim about the moral properties of a certain type of act that is true as a matter of necessity. This is all I need to vindicate the answers that I gave to both the Reductive and the Explanatory Supervenience Objections. In responding to those objections, I drew on the claim that there are at least some moral truths that hold necessarily, and we can reasonably believe that this is the case.

Of course, I may be idiosyncratic in my inability to imagine worlds at which causing T-Pain is obligatory or permissible. The evidence given by what I can imagine is defeasible, for my imaginative capacities might be uncommonly dull. Moreover, our modal intuitions may differ. You may think that you can imagine worlds in which causing T-Pain is morally fine, for example, even if I think that I cannot imagine such worlds. This is a case of modal disagreement – a disagreement over how things might have been.

⁴⁰ To be clear, I'm only making a modest assumption here about the relationship between conceivability and possibility. My assumption is that conceivability is defeasible evidence for possibility. I think that this is plausible, but I don't wish to assume any view about the reasons for its plausibility. For all I know, the strong claim that conceivability entails possibility is true. This view is suggested by Chalmers (1996). Alternatively, it may be that knowledge of possibilities derives from knowledge of essences, and that conceivability only correlates with possibility insofar as it helps us to get knowledge of essences. This view is articulated and defended by Lowe (2012) and Hale (2013). There are other views too, there is a big literature on this topic. But all I need is a modest claim.

Getting into the general problems of modal disagreement among epistemic peers would take us too far afield at present, but there are two things that we can say insofar as it applies to the present discussion. First, one should be careful about how theory affects one's modal judgement. If you already accept a theory that treats morality as mind-dependent, you may find it easy to imagine worlds in which causing T-Pain is permissible (even if you find such worlds abhorrent). But, as I've already said, appeal to morality's mind-dependence cannot be utilised because it has not been established in the dialectic. Equally, an appeal to the mind-independent status of moral reality cannot be utilised, as this has not been established either. It would beg the question if one were to rely on either of these theoretical commitments to vindicate one's modal intuitions about the moral status of causing T-Pain. Even though my intuitions seem to fit with a theoretical commitment to Mind-Independence, I hope that they aren't infected by theory. I don't think that they are, but this is a hard thing to assess.

This brings me to the second point, for it is part of what motivates the modesty of my proposal. I only require that my modal intuitions are reasonable. That is, for the moment my aim is just to show that necessary moral truths are respectable enough to be taken seriously, not that there definitely are such truths. What I have said in this section indicates that intuitions supporting the existence of necessary moral truths are indeed reasonable. This is enough to uphold the Robust Realist's commitment to Non-Reductivism: if the supervenience objections that I have considered fail against any versions of Non-Reductivism that accept necessary moral truths, and if the commitment to necessary moral truths or norms can be reasonably maintained, then Non-Reductivism remains the position to beat.

To clarify, recall from §3.1 that the apparent difference between direction and description constitutes evidence for irreducibility of moral properties. But this evidence is defeasible, and Non-Reductivism is thus vulnerable to being overturned in light of objections. Still, the onus is on critics to provide such objections. And we have now seen that the most influential ways of objecting to Non-Reductivism fail to threaten a version of it that accepts necessary moral truths. Since commitment to necessary moral truths – in the form of moral norms that have modally maximal jurisdiction – can be reasonably held, we have yet to see adequate reason to say that Non-Reductivism has been overturned. Non-Reductivism thus stands.

Of course, Robust Realists do not maintain that *all* moral norms hold as a matter of necessity. Indeed, as I explained in §3.4, some of the more ordinary moral norms are most plausibly seen as contingent. So, if there are both necessary and contingent moral norms, a question arises: how are the necessary and the contingent norms related, given that they have different modalities. I answer this question in some detail in the next chapter, in the course of developing an account of the necessity in Robust Realism. But what matters for now is that, by accepting that there are at least some necessary moral norms, Robust Realists can maintain their external metaphysical Realist Commitment to Non-Reductivism. They can carry on believing that moral properties do not reduce to purely descriptive properties, for the supervenience arguments offered against this view fail.

3.9 Conclusion

The apparent difference between direction and description is evidence for the Robust Realist's commitment to Non-Reductivism. But this sort of evidence is defeasible, and Non-Reductivism is therefore vulnerable to attacks. That is, argument might lead us to reject the appearances. But we have seen that an appeal to the existence of necessary moral norms helps to answer the most powerful ways of objecting to Non-Reductivism. Both the Reductive and Explanatory Supervenience Objections fail to undercut a version of Non-Reductivism that posits necessary moral norms. I have also provided some support to the modest claim that commitment to such truths can be reasonably held. It cannot be quickly dismissed as obviously untenable. So, the crucial lesson of this chapter is that Robust Realists can maintain Non-Reductivism via a commitment to necessary moral truths. In the next chapter, I delve into the nature of this commitment in greater depth.

4 The Necessity in Robust Realism

Saying that Robust Realists should rely on the necessity of (some) moral truths is one thing, but advancing our understanding of these necessary truths is another. My aim in this chapter is to examine the necessity in Robust Realism, in order to sketch a view of how it works. However, the necessity in Robust Realism is a fairly underdeveloped topic. So, whilst I take the claims made here to be plausible, I only make them provisionally, for further inquiry may lead us in different directions. Although there is not time here to consider all viable accounts of the necessity in Robust Realism, I'll contribute by discussing two. I reject one, and provisionally endorse the other. At the end of this chapter I won't say that my proposal is the end of the story, but I will say that we have made headway on the issue. We will have a workable way of understanding of necessary moral truths.

I start with an idea from Terence Cuneo and Russ Shafer-Landau, on which some moral truths are conceptual necessities. I outline this idea in §4.1, and I explain why I am sceptical of it in §4.2 and §4.3. The rest of the chapter develops a different view of the necessity in Robust Realism. This is the view that I'll provisionally endorse, and it starts with the idea that some norms involve 'limiting factors' that restrict their jurisdiction. So, if there are norms that have no limiting factors, they will have modally maximal jurisdiction. I discuss moral norms in §4.4, and limiting factors in §4.5. In §4.6 I show how a commitment to Mind-Independence allows Robust Realists to remove a major limiting factor on moral norms. I thus examine a leading way of defending Mind-Independence in §4.7. I discuss whether moral norms that have had their limiting factors stripped away will be of practical use in §4.8, arguing that some will. I conclude in §4.9 by reiterating the need for further discussion of the necessity in Robust Realism.

4.1 Moral Fixed Points

Cuneo and Shafer-Landau (2014) offer a view of necessary moral truths on which some substantive moral propositions are *conceptual* truths. If this is right, it provides a way to interpret the necessity in Robust Realism as conceptual necessity. This proposal is worth discussing, and if it works it has the potential to take Robust Realism in an interesting new direction. However, I will show that there are important reasons to be sceptical of it.

First, let's consider how Cuneo and Shafer-Landau make sense of the idea that some substantive moral propositions are conceptual truths. Their initial suggestion is that certain non-trivial moral propositions help to determine the boundaries of our moral framework. That is, according to Cuneo and Shafer-Landau, a system of normative claims that does not include these propositions is not a genuinely *moral* system. They name these framing propositions 'moral fixed points,' and propose that they are conceptual truths.¹ The following propositions are suggested as probable moral fixed points:

- It is *pro tanto* wrong to engage in the recreational slaughter of a fellow person;
- There is some moral reason to offer aid to those in distress, if such aid is very easily given and comes at very little expense;
- If acting justly is costless, then, *ceteris paribus*, one should act justly.

We might add 'it is impermissible to deliberately cause T-Pain.' That is, perhaps it is a conceptual truth that it is impermissible to deliberately cause undeserved and undesired pain that is not required for a valuable end.

It might sound strange to say that the necessity in Robust Realism is conceptual necessity. These moral fixed points are substantive moral propositions, after all, but 'conceptual' is often used interchangeably with 'analytic' in this context, and many see analytic truths as strictly formal. On this view, a sentence is true as a matter of analytic necessity if it is true solely in virtue of the meanings of its constituent terms. For instance, the sentence 'all bachelors are unmarried men' is analytically true because 'bachelor' just means 'unmarried man.' Analytic truths are not always this obvious, but they will always be this formal. Moreover, linguistic analysis seems unsuitable for Robust Realism given its commitment to irreducible moral properties. Robust Realists can't allow that basic moral properties and terms analytically reduce in the same way that 'bachelor' analytically reduces.

¹ In addition to their 'framework status,' other factors are said to support treating moral fixed points as conceptual truths. First, if they are true, they are necessarily true. Second, denial of them tends to evoke bewilderment. Third, they are knowable *a priori*. For these claims, see Cuneo and Shafer-Landau (2014: 407-408).

Cuneo and Shafer-Landau seem happy to treat analytic truths as formal and vacuous, and I'm happy to follow them in doing so. But they do not identify conceptual truths with analytic truths. For one thing, they think that conceptual truths can be substantive. In addition to the moral fixed points, Cuneo and Shafer-Landau (2014: 408) suggest the following as respectable candidates for substantive conceptual truth:

- Justified true belief is insufficient for knowledge;
- God's possible existence entails his necessary existence;
- A meaningful statement may be neither empirically verifiable nor analytic.

So, as Cuneo and Shafer-Landau use the terms, conceptual truths can be substantive even though analytic truths are all trivial. And this is one key difference between conceptual and analytic necessity. This is no doubt controversial, but for the sake of argument I am prepared to grant these points.

Another key point is that Cuneo and Shafer-Landau understand analytic necessity, but not conceptual necessity, to be strictly linguistic. Analytic truths are said to be *sentences* that are true solely in virtue of the meanings of their constituent terms. Conceptual truths, however, are said to be *propositions*. Propositions are often expressed through sentences, but not always. As Cuneo and Shafer-Landau note, some propositions "we have yet to discover; others may permanently elude us, owing, perhaps to their complexity" (2014: 411). It might be that all analytic truths are conceptual truths, Cuneo and Shafer-Landau tell us, but it is not the case that all conceptual truths are analytic. So, accepting that the necessity in Robust Realism is conceptual does not require that one also accepts that it is analytic. The so-called moral fixed points are substantive (as well as analytically irreducible) even though they hold as a matter of conceptual necessity.

What is conceptual necessity, then, if it is not analytic necessity? At this point Cuneo and Shafer-Landau (2014: 410) make the important claim that concepts have essences. They suggest, for example, that it is of the essence of the concept 'being human' that it applies to exactly those things that are human. A proposition is conceptually necessary, on this view, if it holds in virtue of the essence of its constituent concepts. More precisely, a proposition [that x is F] is a conceptual truth if it belongs to the essence of ' F ' that, necessarily, anything that satisfies ' x ' also satisfies ' F '

To illustrate how these suggestions are supposed to help us with understanding the necessity in Robust Realism, consider the substantive moral proposition [that it is impermissible to deliberately cause T-Pain]. The idea is that this holds as a matter of conceptual necessity just in case it is of the essence of the concept 'being impermissible' that, necessarily, if anything satisfies the concept 'deliberately causes T-Pain,' then it will also satisfy the concept 'being impermissible.' Building from this thought, Cuneo and Shafer-Landau ultimately end up with the following position:

Moral Conceptual Necessity. There are non-natural moral truths. These truths include the moral fixed points, which are a species of conceptual truth, as they are propositions that are true in virtue of the essences of their constituent concepts.²

As it stands, Moral Conceptual Necessity is compatible with various forms of moral non-naturalism. However, the Robust Realist can embellish it by maintaining that non-natural moral truths are truths about irreducible, non-natural, mind-independent, and categorical moral properties or facts that exist in an external metaphysical sense.³ Moral Conceptual Necessity can thereby provide a neatly fleshed out view of the necessity in Robust Realism.

Cuneo and Shafer-Landau think that Moral Conceptual Necessity has important payoffs, specifically when it comes to answering certain influential objections against views like Robust Realism. They argue that their approach to moral necessities helps us to deal with the existence of persistent moral disagreement, with evolutionary debunking arguments, and with an Explanatory Supervenience Objection of the sort considered in Chapter 3. If this is correct, it will be an appealing approach to take. However, there is reason to doubt Moral Conceptual Necessity. Even if its truth would have the suggested payoffs – something that one might wish to question on another occasion – there are still some important worries about Moral Conceptual Necessity. In the next two sections I develop just one such worry, and in the process I explain why we should be sceptical about taking a moral fixed points approach to the necessity in Robust Realism.

² Cuneo and Shafer-Landau call this 'The Embellished Core Claim,' for it embellishes what they call 'The Core Claim,' which is just the claim that there are non-natural moral truths.

³ As it happens, Cuneo and Shafer-Landau are sympathetic to a robust view of this sort.

4.2 Conceptual Deficiencies

If a proposition is true as a matter of conceptual necessity, then anyone who rejects it can be accused of having made a conceptual mistake. They either fail to possess the relevant concepts, or they possess them but lack an adequate understanding of them, or they possess and understand the relevant concepts but are unable or unwilling (or have otherwise failed) to appreciate what these concepts imply. In short, if we choose to accept Moral Conceptual Necessity we will also have to accept that anyone who rejects the moral fixed points is conceptually deficient in one of these ways.

But we can doubt that all such people are conceptually deficient. Consider the Error Theorist – someone who says (a) that moral discourse presupposes a robust moral reality, and (b) that there is no such moral reality. Error Theorists reject all the moral fixed points, on the basis that accepting them requires what they regard as unacceptable metaphysical commitments. Is the Error Theorist conceptually deficient? Plausibly not. At least, we are entitled to deny that they are deficient in this way. And if we are entitled to deny that Error Theorists make a conceptual mistake in rejecting the moral fixed points, then it is not compulsory to suppose that the moral fixed points are conceptual truths. This, in turn, seriously weakens Moral Conceptual Necessity as a view of the necessity in Robust Realism.

Cuneo and Shafer-Landau are sensitive to an aspect of this worry, for they acknowledge that their proposal might seem *uncharitable* to the Error Theorist. They offer the following remarks as a response:

When we say that error theorists are mistaken ... we are not thereby committed to attributing to them simple-minded positions or flat-footed philosophical mistakes. Rather, our claim is that error theorists are failing to recognize a set of conceptual truths as a result of having been convinced by sophisticated, albeit unsound, philosophical arguments. Assessing the force of such arguments is anything but straightforward ... Those who come down on what we regard as the wrong side of those arguments needn't be making any silly or obvious mistakes (2014: 438).

Even careful thinkers can make conceptual mistakes. The questions that we are discussing are very difficult. We can thus accuse Error Theorists of conceptual deficiency whilst taking them and their arguments entirely seriously.

This seems a decent response to the 'lack of charity' objection, but the matter does not end there. We don't just want to know whether it is uncharitable to accuse Error Theorists of conceptual deficiency. We also want to know whether they are actually conceptually deficient. It is worth stressing that these are different points. To ask if one can accuse Error Theorists of making a conceptual mistake whilst taking them seriously is one thing. To ask if Error Theorists are actually conceptually mistaken is another thing altogether. In saying that Error Theorists are conceptually mistaken, Cuneo and Shafer-Landau are not dismissing Error Theory as silly or stupid. But it is still open to critics of Moral Conceptual Necessity to deny that the Error Theorist makes any conceptual mistake in the first place.

And I suspect that there is a strong case to be made here. Recall that there are a number of ways in which a person might be considered conceptually deficient. First, they might fail to possess relevant concepts. Second, they might possess these concepts whilst failing to satisfactorily understand them. Third, they might possess and understand the relevant concepts whilst failing to appreciate what they imply. It seems clear that defenders of Moral Conceptual Necessity cannot accuse Error Theorists of the first two forms of conceptual deficiency. After all, both parties will largely agree about the nature of moral concepts, and moral discourse in general. Robust Realists and Error Theorists both see moral concepts as presupposing the existence of a robust moral reality, even if they disagree over whether or not this reality actually exists. In other words, the central disagreement between Robust Realists and Error Theorists relates to the existence of moral properties, and not to the best interpretation of moral concepts.

Robust Realists must therefore avoid accusing Error Theorists of failing to possess or understand moral concepts. They might still accuse them of failing to appreciate what these concepts imply, but one can't just assert this. It automatically casts doubt on a theory if it entails that Error Theorists are conceptually deficient, for it is *prima facie* unlikely that they are deficient in this way. So one must make an independent case for the charge to avoid it being *ad hoc*. If the conceptual deficiency charge cannot be validated, we are entitled to deny that Error Theorists are conceptually deficient in rejecting the moral fixed points. This will ensure that it is not compulsory to see the moral fixed points as conceptual truths, which will undercut Moral Conceptual Necessity as a view of the necessity in Robust Realism.

We thus need to consider how one might support the claim that Error Theorists are conceptually deficient. If no support for this claim can be found, we should be sceptical about the tenability of Moral Conceptual Necessity. In the next section, I draw out a case for a conceptual deficiency charge from Cuneo and Shafer-Landau's discussion, before arguing that it fails.

4.3 Against Moral Fixed Points

How might defenders of Moral Conceptual Necessity support the charge that Error Theorists fail to appreciate what their moral concepts imply? An idea that we might glean from Cuneo and Shafer-Landau's discussion is that the Error Theorist is misled by a certain methodology. For it is a mistake, according to Cuneo and Shafer-Landau, to reject highly evident moral propositions (such as the moral fixed points) simply on the basis of certain highly controversial metaethical claims (such as the claim that our moral beliefs can be debunked, or that moral reality is unacceptably 'queer').

Cuneo and Shafer-Landau diagnose Error Theorists as falling foul of this "suspect philosophical methodology" (2014: 438). What exactly they are getting at here is not fully clear, but maybe they are gesturing to something like the following line of thought. In philosophical inquiry, best practice generally involves seeking to accommodate our intuitions about our chosen topic by developing theories that are able to vindicate these intuitions. Error Theorists thus fail to follow best philosophical practice, for their view repudiates our strongest moral intuitions by rejecting the moral fixed points. The idea is thus that we can provide indirect support to the accusation of conceptual deficiency by pointing to a questionable methodology – Error Theorists miss what their concepts imply due to a distorting process of philosophical reasoning, one that gives too much weight to contentious metaethical claims, and too little weight to moral intuition.

If this is true, it means that the conceptual deficiency charge is not just an *ad hoc* claim that one is forced to make if one already accepts that some substantive moral propositions are conceptual truths, for the idea that Error Theorists make a conceptual error is independently supported by a diagnosis of their methodological situation. However, if this is what Cuneo and Shafer-Landau are getting at, it is unconvincing. There are two problems with trying to support the conceptual deficiency charge in this way.

First, the methodology of which Cuneo and Shafer-Landau are suspicious is widely employed in philosophy. Claims that seem intuitively plausible are very often given up in light of deeper reflection and more considered argument. To illustrate, intuitively it seems highly plausible that time is real, that scientific investigation delivers correct theories of reality, that human beings have free will, and that most moral agents have reasonably stable character traits. Yet all of these claims are subject to apparently legitimate debate, and there are philosophers who reject them by appeal to contentious philosophical (and empirical) arguments. If this methodology is suspect when it is used by the Error Theorist, then it is presumably suspect when used by any philosopher. Perhaps it is indeed a bad methodology wherever it is used, but we must recognize that a wide range of fields in philosophy would be radically altered if philosophers were directed only to construct theories that accommodate our strongest intuitions.

Second, it is not actually clear that Error Theorists operate with a bad methodology. To see this, consider the following as a rough guide to philosophical inquiry. *Step One*: sort out your intuitions. You may find that some of them seem to conflict, so you will need to deal with that. On closer examination, it might be that the conflicts are merely apparent. If they're not, however, you'll have to repudiate some intuitions. You should give priority to those intuitions that seem to you most credible. You are now ready to proceed to *Step Two*. *Step Two*: develop a theory that can vindicate your considered intuitions, and test this theory for its intuitive acceptability, its internal coherence, its consistency with the known facts, and so on. If it passes these tests, accept the theory. If it doesn't, move to *Step Three*. *Step Three*: try revising your theory, and perhaps some of the least evident intuitions that the theory has been designed to support. If the theory remains problematic despite your efforts to get it to work, you may find that you need to move to *Step Four*. *Step Four*: accept that the intuitions with which you started are in error. Explain why this is so. You might also explain why they seemed so plausible, even though they are in error.

Although it is stated a little roughly, I take it that this four-step programme is a familiar process of philosophical theorising. It also seems obvious that this methodology is a legitimate one. So, what matters here is whether we can see the Error Theorist as someone who has taken this four-step programme with respect to the metaethical domain. And we can.

Most people are likely to find that their moral intuitions, including the moral fixed points, seem highly evident. So (like most people) the Error Theorist begins with Step One by getting their moral intuitions in order. They then proceed to Step Two, and (like the Robust Realist) they judge that they need to invoke a robustly realistic metaethical theory in order to vindicate their moral intuitions. However, in testing this theory, the Error Theorist (unlike the Robust Realist) concludes that it is in some way unacceptable. They might judge that its metaphysical commitments are 'queer,' for example. So they move on to Step Three, at which point they attempt to revise their robust theory in such a way that it avoids any awkward metaphysical commitments. But now the Error Theorist finds that this revised (non-robust) theory in some way fails to vindicate their considered moral intuitions. Despite their best efforts, they can't find an appropriate balance between their moral intuitions and a theoretically acceptable account of them. In the end they judge that they have to move to Step Four, at which they acknowledge and explain the error underlying those apparently obvious moral intuitions. In short, they arrive at Error Theory.

Now, we can debate whether the Error Theorist is right to think that a robust moral metaphysics is unacceptable. That is, one might say that the Error Theorist makes a mistake at Step Two; this is what Robust Realists can say. Alternatively, one might suggest that the Error Theorist makes a mistake in thinking that a non-robust theory cannot vindicate at least a respectable chunk of our moral intuitions. That is, one might say that the Error Theorist makes some mistake at Step Three; this is what defenders of non-robust theories can say. Either way, there doesn't seem to be anything wrong with the Error Theorist's *methodology*. Indeed, it seems that they have proceeded entirely legitimately along our plausible four-step programme. The significance of this is that it undermines the support that we were hoping to see provided for a conceptual deficiency charge.

Cuneo and Shafer-Landau seemed to suggest that the conceptual deficiency charge can be supported by a diagnosis of the Error Theorist's methodological situation, but the diagnosis in question is implausible. This means that defenders of Moral Conceptual Necessity cannot appeal to it as an independent justification for the claim that Error Theorists are conceptually deficient in rejecting the moral fixed points. The conceptual deficiency charge is thus unsupported, so it is not compulsory to accept it.

If we are entitled to deny that Error Theorists are conceptually mistaken in their rejection of the moral fixed points, then we do not need to accept that the moral fixed points are conceptual truths. We must thus remain sceptical about the idea that the necessity in Robust Realism is conceptual necessity. Of course, defenders of Moral Conceptual Necessity might seek some other justification for the conceptual deficiency charge. Even if the Error Theorist makes a mistake at Step Two or Step Three of the four-step programme outlined above, it is still possible that this is a conceptual mistake. The defender of Moral Conceptual Necessity thus has to engage in the project of finding an independent way of supporting the claim that Error Theorists make some conceptual mistake at one of these points. I am pessimistic about this project, however, and I think that we are better off looking elsewhere for an account of the necessity in Robust Realism.⁴

4.4 Moral Norms

We saw in Chapter 3 that necessary moral norms help the Robust Realist with the Explanatory Supervenience Objection. So perhaps an account of the necessity in Robust Realism can begin with the idea of a moral norm. Norms abound. We often call them rules, standards, prescriptions, laws, principles, guidelines, constraints, codes, strictures, and so forth, but in general their role is to specify directions for agents within given sets of circumstances. In other words, a norm governs what counts as legitimate behaviour in a set of relevantly similar situations, and one can thus satisfy or fail to satisfy a norm depending on how one behaves. If a norm directs humans to exercise three times a week, then I – a human – satisfy it by exercising three times a week, and fail to satisfy it by exercising twice a week.

The norms that govern legitimate behaviour in a situation may be implicit, but in principle they can be stated in directive sentences using terms like ‘should’ and ‘ought’ and ‘must.’ So, norms are satisfaction-apt. As we saw in Chapter 3, it is satisfaction-aptness that makes it hard to see how the moral could be captured in purely descriptive terms. A sentence like ‘deliberately causing T-Pain is impermissible’ has the same form as ‘the table is brown,’ but the former has a directive aspect that the latter lacks. Such claims involve or imply direction that one can satisfy or fail to satisfy.

⁴ A shorter version of §4.1 to §4.3 has been published as Ingram (2015b).

In addition to moral norms, there are epistemic norms, prudential norms, legal norms, professional norms, and so on. There is a question to ask about the relations among types of norm – maybe professional norms are a type of prudential norm, for example. The boundaries can blur, and they are often hard to pin down, but the distinctions are useful. There is moral region of the broad normative domain, though it may well overlap with other aspects of this domain. Various criteria have been suggested to demarcate this specifically moral region. As we'll see in Chapter 5, one is that moral norms play a role in interpersonal justification with respect to conduct. Another is that they are norms that tend, when we obey them, to help us achieve *eudaimonia*. Yet another is that they are norms whose satisfaction realises the good, where the good is indefinable but knowable via intuition. I don't know exactly what makes a norm moral, but I assume that there is a category of norm that we can usefully classify under this term.

At first glance, norms seem like a distinct type of entity. The idea of a norm is not that of a substance, fact, property, relation, or kind. They appear to form a distinct category of being. Having said that, norms do seem a bit like *laws*, for both govern behaviour. Laws of nature govern the behaviour of entities studied by the natural sciences. Moral norms are not laws of nature, but they do govern the behaviour of moral agents. They may fail to govern *actual* behaviour – it is possible to disobey them – but norms govern us by fixing whether our conduct counts as legitimate. So, even if norms are not laws of nature, we can see them as a special type of law.

Are any norms explicable via a more basic category of being? Yes, in some cases, for many norms are purely social entities. Whilst it is often pointless to explain social categories in more basic terms, it is typically possible in principle. For instance, one may think that professional norms simply reflect descriptive facts about how we have agreed to behave in a certain context. Such norms are not part of a distinctive category, for they can be re-categorised as descriptive facts. However, other types of norm won't submit to a re-categorisation of this sort, for they are irreducible. They are not a subset of the category substance, or property, or fact, or whatever. These norms appear to form a distinct category of being. Thus, whilst even the most fundamental professional or prudential norms may be re-categorised at a more basic level of explanation as facts, the most fundamental moral norms form their own category, as a distinct type of entity.

Before moving on, note that I defined Robust Realism (in §1.1) as a view about moral properties and not moral norms. Given my current emphasis on the role of norms, it is worth stating why I defined it in this way. I did so because, though I argued in Chapter 3 that Robust Realists should give norms a key role in their metaphysics, we must allow that one can be a Robust Realist without appealing to norms at all. Most obviously, a ‘particularist’ – someone who is hostile to the idea of general moral norms – can be a Robust Realist.⁵ That is, a particularist can endorse the Realist Commitments whilst viewing them in an external metaphysical sense, and will thereby be a defender of Robust Realism. It’s simply that this is not (I claim) the best form of Robust Realism, for the best form of Robust Realism involves necessary moral norms. As I said in Chapter 3, such norms help us with the Explanatory Supervenience Objection. In the next section, I begin to develop an account of moral necessity by looking at how contingent factors limit the range of many non-fundamental moral norms.

4.5 Limiting Factors

A moral norm, by its nature, has jurisdiction across some set of worlds. It won’t just apply in a one-off circumstance, for moral norms universalise. But we can ask about the modal range of a moral norm. A norm may have authority in some worlds but not others, or it may be authoritative in all worlds. What the Robust Realist needs is an understanding of how at least some moral norms are authoritative in all possible worlds. But let’s start with the fact that, even for the Robust Realist, not all moral norms are necessary. Consider a moral norm that prohibits recreational violence. This is a norm that plausibly has authority in our world, and in worlds like ours. But it doesn’t have authority in all worlds. For instance, there is a world whose inhabitants can only experience pleasure whilst others are being recreationally violent toward them. This is a peculiar but possible world, in which recreational violence is the only means to a valuable end. Plausibly, the norm prohibiting recreational violence has no authority in this world. So, even if there are necessary moral norms, this one is clearly contingent. I propose that we might begin to understand the necessity in Robust Realism by providing an explanation of these contingent parts of morality.

⁵ For more on particularism, see Dancy (1993; 2004a) and Little (2000).

We can begin to see how a norm prohibiting recreational violence gets to be authoritative in some worlds but not others by noting that it is non-fundamental. That is, this norm in some way rests on or presupposes a deeper, more fundamental norm. For instance, the norm that prohibits recreational violence plausibly presupposes a deeper norm according to which it is wrong to cause undeserved and undesired pain that is not required for any valuable end – that is, a norm that prohibits deliberately causing T-Pain. This means that a norm prohibiting recreational violence has authority in our world because (a) it is impermissible to deliberately cause T-Pain and (b) to participate in recreational violence in our world is almost always to cause T-Pain. But there are worlds in which engaging in recreational violence never causes T-Pain, for instance worlds in which the only available means to a certain valuable end, like the experience of pleasure. A norm prohibiting recreational violence lacks authority in such worlds.

To clarify, in our world, a norm prohibiting recreational violence has authority, and this involves two things being the case. First, that there is a more basic norm that prohibits deliberately causing T-Pain. Second, that to engage in recreational violence in our world is almost always to cause T-Pain. But this second aspect of the norm's authority in our world is contingent, for there are worlds in which participating in recreational violence is the only way to cause pleasure. Recreational violence will thus be permissible (maybe even obligatory) in such worlds, as the only means to that valuable end. The fact that it relies upon this second element thus ensures that a norm prohibiting recreational violence only has contingent authority.

Put another way, the norm is contingent because it presupposes a contingent fact. This contingency limits the worlds at which the norm has authority. It is a 'limiting factor' on the norm's modal jurisdiction. Of course, this is not to deny that the norm has authority in many worlds. But its authority comes from the fact that it presupposes a more basic and perhaps necessary moral norm, *in combination with* a contingent fact that is actualised within some worlds. In our world, engaging in recreational violence tends to cause T-Pain, so a norm that prohibits it has authority here. But there are also worlds in which engaging in recreational violence never causes T-Pain. In such worlds, a norm prohibiting it lacks authority. This is the source of its contingency. The contingent presuppositions that enable this norm to apply at some worlds are also why it fails to apply at others.

More abstractly, a contingent moral norm *inherits* its contingency from a contingent limiting factor that is in some way presupposed by it, even though it also presupposes a more basic moral norm that may count as necessary. The authority of the more basic norm, combined with the specifics of the limiting factors, determine the range of worlds at which the non-fundamental norm has authority. Now, this suggests that there is a hierarchy of norms. I don't know if there is a fundamental norm that can unify all others but, at some point, the hierarchy of norms will end with either one basic norm, or a plurality of equally basic norms. Either way, it is the most basic moral norms are the clearest candidates for necessity. They don't rest on any limiting factor, for they don't rest on anything. In other words, they don't involve any contingent fact capable of restricting their modal jurisdiction. They are thus 'unlimited,' and have what David Enoch calls 'modally maximal jurisdiction.' They have authority in every world, and this gives us a way to model the necessity in Robust Realism. The necessary moral norms are those that presuppose no limiting factors. These include (but are perhaps not limited to) the fundamental moral norms.⁶

We can clarify the suggested way of interpreting the necessity in Robust Realism by looking at an analogous but non-normative case from mathematical domain. Consider a straightforwardly mathematical claim like '19 is an odd number.' This, we can agree, is a necessary truth. Maybe the necessity of this mathematical truth involves or presupposes some other necessary mathematical claim that is more fundamental than it, for instance the necessary mathematical truth that a number is prime if and only if it is a natural number greater than 1 that has no positive divisors other than 1 and itself. But it doesn't appear to involve or presuppose any contingent element. It is thus modally unlimited, and true at all worlds. Now consider the claim '19 is an odd number, though I wish it was even.' This is true, but contingently. It is contingent despite the fact that it is partly composed of a necessary mathematical truth. This claim is *made* contingent by its inclusion of my wish, which I don't have at all possible worlds.

⁶ I'm suggesting here that all of the fundamental moral norms are necessary, but I'm not suggesting that all of the necessary moral norms are fundamental. Many non-fundamental norms are contingent, but some norms might be necessary despite presupposing some more basic necessary norm. For instance, in my view a norm that prohibits causing T-Pain is a candidate for necessity, as it doesn't seem to involve any limiting factor. However, we should be open to the possibility that it rests on a more basic necessary norm.

Put another way, the statement inherits its contingency from this contingent bit of its content. Without the contingent ingredient, it would be necessary. Adding my wish limits the range of worlds at which the statement is true. It is a 'limiting factor,' as I am putting it. And, as we've seen, there are also limiting factors in the case of many non-fundamental norms. A claim like 'it is impermissible to engage in recreational violence' only has its authority insofar as it presupposes the fact that recreational violence almost always causes T-Pain.⁷ But there are possible worlds in which it doesn't do this. The link between recreational violence and the deliberate causing of T-Pain is contingent. And this contingency is what constitutes the limiting factor on the modal jurisdiction of the moral norm that prohibits recreational violence. So, if there moral norms that do not involve (or presuppose) any contingent elements – that is, any limiting factors – then they will have unlimited authority. They'll have a modally maximal jurisdiction. In short, their authority will extend to all possible worlds.

In sum, certain moral norms are contingent because they involve or presuppose at least one limiting factor – that is, at least one contingent element that restricts the range of possible worlds over which the norm has authority.⁸ To establish whether there are necessary moral norms, we thus need to consider whether there can be norms do not involve or presuppose any limiting factors. That is, Robust Realists can understand a moral norm as holding necessarily if it is not limited by any contingency in its content or in its presuppositions. These unlimited moral norms, if there in fact are any, provide a way to understand the necessity in Robust Realism.

⁷ If this remains unclear, it may help to consider the cancellability of the presupposition. For instance, as I have suggested, 'it is impermissible to engage in recreational violence' presupposes that recreational violence is almost never the means to some valuable end. We can cancel this by saying 'it is impermissible to engage in recreational violence, unless it is the only means to a valuable end.' When we cancel the presupposition we indicate that there are contexts in which the norm lacks authority. Specifically, contexts in which engaging in recreational violence is the only means to a more valuable end. Since it is clear that such contexts are possible, there are worlds in which the norm lacks authority. Thus, it is contingent due to its presupposing a contingent limiting factor.

⁸ Cf. Scanlon (2014), who distinguishes pure normative claims, which don't involve any non-normative claims – non-normative bits are "subjunctivised away" (2014: 40) – from mixed normative claims, which involve both a pure normative claim and some contingent non-normative claims. This distinction, also found in Schroeder (2014), is similar to what I have in mind, but I have not drawn on the idea of any bits being 'subjunctivised away.'

To get clearer on the idea of a limiting factor, and on how we can go about trying to strip away such factors to discover the necessary or modally unlimited moral norms, I'll consider an important case study in §4.6.

4.6 A Case Study: Mind-Independence

Robust Realists take moral reality to be objective in the following strong sense:

Mind-Independence. Moral properties are constitutively independent of any agent's or set of agents' actual or hypothetical attitudes, beliefs, and conventions.

Although this claim is about moral properties, Robust Realists can expand it to cover the mind-independence of moral norms. That is, they can say that moral norms govern us independently of our attitudes, beliefs, and conventions. In this section I show how Mind-Independence, interpreted as an external metaphysical claim, strips away a few big limiting factors. Now, to eliminate dependence on attitudes, beliefs, and conventions is to go some way toward necessity. Not all the way, by any means, for these are not the only contingent elements that may limit the range of worlds across which a norm has its authority. But we can use the commitment to Mind-Independence to illustrate the process of stripping away limiting factors that will help Robust Realists to identify robustly necessary moral norms.

We can begin to make the point by observing that at least certain simple mind-dependent views of moral reality involve significant limiting factors. Consider, for example, the following metaethical theory:

Metaethical Conventionalism. The moral norms that govern human communities are conventional norms, which reduce to descriptive facts about agreements that hold within those communities.

This is a reductive view, for it takes moral norms and reduces them to the conventions of a community. It then reduces these conventions to certain descriptive facts about the agreements that hold among members of that community. Here we can allow that these agreements may be implicit or explicit, and we can allow there may be members of the community who dislike the relevant conventions, and who will get rid of them if they can. What matters is that have agreed (tacitly or otherwise) to be governed by them.

Metaethical Conventionalism is to be interpreted as a view about the nature of morality's normativity. It tells us that the authority of moral norms is ultimately a matter of the agreements that hold within a given community. As it is stated here it is clearly a rather simplistic metaethical theory, not one that many metaethicists are likely to defend. But that's ok. I'm not interested in the plausibility of Metaethical Conventionalism, for I just want to use it to illustrate how certain mind-dependent theories of moral norms are (modally speaking) inherently limited.⁹ We can then see how a mind-independent metaethical theory will eliminate these limiting factors.

Metaethical Conventionalism can be usefully compared to certain other ways in which conventions seem to bind us. Consider the following:

The Rules of Tennis. The rules that govern the community of tennis players are conventional norms, which reduce to descriptive facts about agreements that hold within that community.

Plausibly, the rules of tennis have a certain sort of authority over tennis players. For instance, by the rules of tennis, before one serves one must stand behind the baseline without touching it. So, if I serve with my foot on or over the baseline – that is, if I break the 'foot fault' rule – the umpire can truly say to me 'you should not have done that.' This is a normative claim, and it has a certain type of authority. And, as *The Rules of Tennis* says, it has this type of authority because the members of the community agree to be governed in certain ways. More broadly, the authority of the rules of tennis is said to be a matter of certain descriptive facts concerning an either implicit or explicit agreement within the community of tennis players.

The rules of tennis are not sewn into the fabric of reality. They bind in a meaningful sense, but their authority is mind-dependent. This is because it rests on the members of the relevant community having agreed to obey certain rules. Metaethical Conventionalism makes a similar claim about morality. Of course, we tend to care far more about morality than tennis. However, if Metaethical Conventionalism is correct, the authority of moral norms is of the same basic sort as the authority of the rules of tennis.

⁹ More sophisticated views employ rigidifying manoeuvres to extend the range of worlds across which moral norms have jurisdiction. I set this aside, for I'm just using this simple view to bring out a point about robust theories, not to criticise non-robust theories.

Now, the mind-dependence of the rules of tennis limits the range of worlds at which these rules have authority. This is because there are worlds at which the agreements that establish the authority of these rules are not in place. There are worlds, for example, in which a community has agreed to follow the rules of tennis *except for the foot fault rule*. In this world, if I were to break this rule by serving with my foot on the baseline, the umpire couldn't correctly say 'you should not have done that,' for the convention that this presupposes is not actualised. The foot fault rule thus has authority in some worlds but not others, because the agreements that establish its authority are contingent – they exist at some worlds and not others.¹⁰

In other words, the fact that the foot fault rule is mind-dependent ensures that it is merely contingently authoritative. Its authority depends on our agreements, and these agreements exist in some worlds but not in others. So, the mind-dependence of the rules of tennis is a limiting factor on their modal jurisdiction. They have a certain sort of authority, but it is contingent. As you might guess, the same point can be applied when we come to Metaethical Conventionalism. The social conventions that are said to govern our moral community have authority in a restricted set of worlds. Specifically, worlds at which the same agreements are in place among members of the moral community. And, since there are worlds in which these agreements are not in place, moral norms are contingently authoritative. At least, this is the case if Metaethical Conventionalism is true.

Of course, if the authority of moral norms were not just grounded in agreements, things might be different. In particular, if they were taken to direct us independently of our attitudes, beliefs, and conventions, then the range of worlds at which they have authority would be less restricted. This simple mind-dependent view therefore involves limiting factors that mind-independent views do not. By committing to Mind-Independence, the Robust Realist thereby ensures that the set of worlds at which moral norms have authority is not limited by our contingent agreements. Nor, indeed, by contingent attitudes or beliefs. It may be limited in other ways, but Mind-Independence alone takes us part of the way to necessary moral norms.

¹⁰ Some rules may be constitutive of tennis; one cannot alter them without making a new game. The foot fault rule is not one. If you disagree, pick another rule. There must be some contingent rules of tennis. Its rules do change, and it isn't a new game with every change.

Mind-Independence is thus a useful case study in the elimination of limiting factors that, on the view proposed here, helps us to understand the necessity in Robust Realism. Whilst certain mind-dependent theories suggest that moral norms can only possess a certain sort of authority by involving or presupposing contingent attitudes, beliefs or conventions, by treating moral norms as mind-independent the Robust Realist removes such contingencies. By seeing moral norms as constitutively independent of our attitudes, beliefs, and conventions, the Robust Realist strips away several factors that would limit the range of worlds at which moral norms have their authority. A commitment to Mind-Independence is thus part of the process of eliminating limiting factors, taking us some (but not all) of the way toward moral norms whose jurisdiction is unlimited or modally maximal.

Note that, like every reductive view, Metaethical Conventionalism has first-order implications. If we accept its second-order claim about the nature of moral normativity, then we must accept the first-order claim that one ought to obey the community's conventions. However, we're not dealing here with a *purely* first-order view. After all, there can be a purely first-order conventionalism – even Robust Realists can take a first-order view on which one must always follow the community's conventions. This purely first-order position does not count as a rival to Robust Realism, for it can be coherently combined with it. It thus matters that Metaethical Conventionalism is not just a first-order view. It is a second-order view that has some first-order implications, and is therefore a rival to Robust Realism.

Hang on. If a Robust Realist adopts a first-order conventionalism, won't their view of moral norms be just as modally limited as that of the Metaethical Conventionalist? No. Whilst a robust conventionalist has to hold that we must obey the community's conventions, this 'must' is meant to be necessary, or mind-independent. The conventions themselves are contingent on agreements, but for robust conventionalists the claim that one must follow them is not. In contrast, the Metaethical Conventionalist cannot see the 'must' as necessary without altering their view. They may even have to accept Mind-Independence, to remove the contingency that limits the jurisdiction of the 'must.' Without some revision to their view, a Metaethical Conventionalist can only see the claim that one 'must' obey a community's conventions as authoritative if they take it to presuppose an agreement to do just that. And this will bring in a contingent limiting factor.

In sum, the Robust Realist's commitment to Mind-Independence is a useful illustration of the process involved in stripping away limiting factors. It helps to see how we can move toward necessary or unlimited moral norms. Mind-Independence doesn't take us all of the way, but by removing dependence on contingent attitudes, beliefs, and conventions, it does take us some of the way toward the necessity in Robust Realism. Because it is an important to the necessity in Robust Realism, and because it is anyway an important part of Robust Realism itself, in the next section I briefly discuss a common way of trying to motivate and defend the idea that there are mind-independent moral norms. As we'll see there, and in Chapter 6 as well, this common defence of Mind-Independence is in some ways unsatisfying. But it is interesting, and it is worth exploring in more depth.

4.7 More on Mind-Independence

After all, suggesting that the necessity in Robust Realism is to some extent tied to the commitment to Mind-Independence is one thing. Defending that commitment is another. In this section, I will outline a common way of arguing that moral properties or norms are constitutively independent of our attitudes. I will show in Chapter 6 that this way of arguing about moral reality fails to be dialectically effective, but it is worth introducing it here in order to clarify this element in our account of necessary moral norms.

The way of arguing in question draws on intuitions about certain cases – cases in which someone acts in a way that is intuitively wrong, but in which there is no attitude or convention (or whatever) that we can use to make sense of that intuition. Metaethical Conventionalism will help to clarify this way of arguing. On that view, moral norms get their authority from the members of a community having agreed (implicitly or explicitly) to be ruled by certain conventions. To show that this is false, defenders of Mind-Independence might ask us to imagine a community where a moral norm about which we are particularly confident is not backed up by any relevant convention or agreement. They then suggest that one ought to act as that norm directs, even if one cannot appeal to any convention or agreement to make sense of this claim. The basic thought is that we have substantive normative intuitions about the correct moral norms, and that mind-dependent theories like Metaethical Conventionalism are unable to capture all of these. This is taken to support a mind-independent view of morality.

We can clarify this style of argument by imagining a community with a set of moral conventions that direct agents to perform an annual sacrifice of one teenager, selected by a lottery system. The thought is that we want to be able to say, not just that we personally *dislike* the norms of that community, but that its members make a *mistake* about which norms are correct. But we have no resources to make that sort of claim if we hold that the authority of moral norms derives from conventional agreements and nothing else. All we can say is that, whilst the sacrifice of a teenager is genuinely obligatory given the conventions of the given community, we dislike those conventions, and will make different agreements when we decide how our community should be run. However, many consider this to be unsatisfactory. They say that, in order to vindicate the intuition that one ought not to sacrifice teenagers, we have to postulate a moral norm that prohibits such behaviour even if there are no agreements to back it up.

At this point the Metaethical Conventionalist (or indeed any other defender of a mind-dependent view, for the point is meant to generalise) might reply by suggesting that they do have resources to criticise norms prescribing the sacrifice of teenagers. First, they might propose that the community accepting such norms is in some way incoherent – maybe the conventions that prescribe the sacrifice of teenagers are in conflict with other conventional norms accepted within the relevant community, like conventional norms that forbid causing harm to innocent members of the community. In that case, we can accuse the community of being irrational in the sense that their conventions are inconsistent and thus in need of revision.

Second, the Metaethical Conventionalist may say that the relevant convention derives from false non-moral beliefs. Even if the community is consistent in its conventions, it may have agreed to a sacrificial practice because of some falsity. For instance, they may have based the agreement on the need to please a certain higher being and, if no such higher being exists, the norm can be thus criticised and rejected because it rests on a false non-moral belief. In short, those who understand moral norms as mind-dependent can appeal to requirements of both coherence and true non-moral belief in order to capture the intuition that those who annually sacrifice a teenager ought not to do so. No mind-independent norms are needed. All we need to do is show that those who practice sacrifice work from an incoherent set of norms, or norms that rest on false non-moral beliefs.

However, defenders of Mind-Independence will not be convinced that this is sufficient to deal with the problem. After all, they can point to the possibility of a community that prescribes the annual sacrifice of a teenager whilst having coherent conventions and true non-moral beliefs. If this community has no norm that forbids causing harm to innocents, there may be no conflict in their conventions. You might think that such a community will be unstable, and that it won't last long without forbidding the harming of innocents. But perhaps it can rub along quite nicely with a more fine-grained norm that forbids harming innocents *unless they are teenagers who have been selected by lottery for the annual sacrifice*. This norm can stabilise the community and doesn't conflict with the norm that requires the annual sacrifice of teenagers. Moreover, the community may engage in this practice of sacrificing teenagers not to please or appease a non-existent higher being, but just because it gives them pleasure. So, it is consistent in its conventions and agreements, and it doesn't rest these strange conventions on false non-moral beliefs. Such a community is odd, and maybe no community like it has existed in the actual world. Still, it is possible.¹¹

The idea is that Metaethical Conventionalism, like all metaethical theories, should give the right answer even in odd or far-fetched cases. If it fails to do so, that seems to be a problem for the theory. At least, that's the idea that is commonly put by defenders of Mind-Independence, and it generalises to pretty much all mind-dependent views, for all such views presumably have to allow the possibility of coherent and well-informed individuals and communities with very different attitudes to our own.¹² In such communities, different norms and values and reasons will obtain. Those who defend Mind-Independence will thus accuse mind-dependent views of giving incorrect answers to questions about what is right and wrong in certain cases, and often argue from this to Mind-Independence itself.

¹¹ The community imagined here is akin to a type of character who Street (2009) names the 'Ideally Coherent Eccentric.' I discuss this character further in Chapter 6.

¹² Maybe it doesn't generalise to views that take on at least some Kantian commitments, like the theories developed by Smith (1994) and Korsgaard (1996; 2009). For instance, Korsgaard says that moral norms are self-legislated, but that the nature of rational agency ensures that we are all rationally committed to self-legislating some moral norms and not others. So this is an ultimately mind-dependent view, but it is meant to accommodate the sort of intuition to which defenders of Mind-Independence are appealing. I am sceptical about these Kantian views, but I set them aside to be considered on another occasion.

I should note, however, that I am not keen on this way of arguing. Arguments that rely on substantive moral intuitions are unlikely to lead to metaethical progress. I began to explain why at the end of Chapter 2, and I say more on it in Chapter 6. I raise it here because, by eliminating certain limiting factors, Mind-Independence is key to robustly necessary moral truths. If this common way of defending it is problematic, then the Robust Realist has reason to worry. So this, if you like, is an initial warning that we need to more carefully evaluate the tenability of a commitment to Mind-Independence. That's the job of Chapter 6, however. In the next section, I will ask a different question about the nature of necessary moral norms.

4.8 Necessary Norms in Ordinary Life

What role do necessary moral norms, interpreted as norms that lack any contingent limiting factors, play in ordinary moral life? One possibility is that they don't have an interesting role to play here, for it may be that the process of stripping away limiting factors will leave us with norms that are staggeringly complex. Such norms might be too unwieldy to do useful work in ordinary life. Even if we can use such entities to understand the necessity in Robust Realism, one might worry that they do not deserve to be called 'norms.' My suspicion is that it doesn't ultimately matter if there is no significant role for necessary moral truths in ordinary moral thought and discourse, but it would be nice if can find some sort of role for them. I think that we can find such a role, and in this section I will explain what it is.

Let's start with a useful distinction made by Sean McKeever and Michael Ridge in a discussion of particularism.¹³ The initial distinction is between norms as *standards* and norms as *guides*. Standards are "entirely exceptionless generalisations" and "should be understood as necessary moral truths" (2006: 7). However, whilst standards apply in all possible worlds, they might be useless in practice. Perhaps so much so that we will be reluctant to call them genuine norms. Specifically, a standard might be so complicated that human beings will find them impossible to apply and internalise. These standards are of little use in everyday life. So, standards are necessary or unlimited moral truths, but they can fail to give practical guidance.

¹³ I haven't engaged with particularism directly, but I noted earlier that I reject it. So do McKeever and Ridge, and I am sympathetic to many of their responses to it.

As you might expect, this is not a problem with norms that are understood as guides. McKeever and Ridge see guides as everyday norms that “provide useful direction to a conscientious moral agent,” where this usefulness involves something like “reliably leading the agent to perform the right action for the right reasons” (2006: 8). Guides are just rules of thumb, like ‘it is wrong to lie’ and ‘one ought to be kind.’ These norms are helpful in everyday life, for in most cases they get us to do the right thing for the right reasons. But they are too rough to be necessary truths. We can easily imagine cases in which lying or being unkind are the right ways to go. Guides are useful, then, but are too rough to eliminate all contingent limiting factors. So, whilst standards are modally unlimited but very often practically unhelpful, guides are modally limited but of great practical use.

What we need are norms that combine the necessity of standards with the practicality of guides. We need *action-guiding standards*. There might be a number of ways of developing this notion. McKeever and Ridge (2006: 10) offer one that seems pretty promising. They propose that an action-guiding standard is a standard that necessarily plays a role in the moral psychology of an ideally virtuous moral agent. It may not figure in such an agent’s conscious moral reflection, at least in ordinary cases, but it plays some role in their moral psychology. For one thing, it is there as the final arbiter of what one ought to do. When everyday rules conflict, one may appeal to an action-guiding standard to settle what is right in the specific case. As a rule of thumb, it is wrong to lie. And as a rule of thumb, it is obligatory to be kind. But how should I act if lying is the kindest thing to do? I can answer this question by appealing to a deeper action-guiding standard.

So, rules of thumb presuppose action-guiding standards. The rule that prohibits lying and the rule that demands kindness rest on deeper moral norms that, in combination with contingent facts about our world, determine their authoritativeness in many situations. This fits with what I said in §4.5; a contingent norm that prohibits recreational violence rests on a necessary norm that prohibits causing T-Pain. What matters here is that part of the point of the contingent norms is that they are useful rules of thumb. A norm that prohibits causing T-Pain is not wildly complex, but it involves elements that can be ignored in many situations. It is part of the psychology of virtuous agents, insofar as it grounds the rules of thumb that guide such agents, but it is only brought to the surface in moments of conflict.

I take it that any substantive theory can make use of the notion of an action-guiding standard, but the easiest way to see how they might be used is to look to the ‘two-level’ utilitarian theories of Henry Sidgwick and R.M. Hare.¹⁴ On this view, the foundational utility norm is rarely (if ever) employed in our ordinary moral thought. For instance, if you must choose between (a) keeping your promise to help your friend to move house and (b) taking the dog to the vet for a relatively minor injury, you don’t think of the norm that directs one to maximise utility. You consider the duty to keep one’s promises, the significance of doggy welfare, and so forth. But, according to two-level utilitarians, such considerations count as duties or have significance ultimately in virtue of the foundational utility norm. So, the notion of action-guiding standard plays an indirect role in your moral life, by determining the rules of thumb by which you live, and perhaps by settling the conflicts that sometimes emerge when you apply the rules of thumb.

Whilst the Robust Realists may not *need* necessary moral truths or norms to play a role in ordinary life, there is thus a plausible way in which they might do so. If there are action-guiding standards – that is, necessary norms that play some role in the moral psychology of an ideally virtuous agent – then there is a way in which necessary moral norms have practical significance. Of course, one might object that there are no such norms. In other words, one might think that the only way for a norm to be necessary is for it to include a near infinite disjunction of exceptions, a huge list of limiting factors being stripped away. In that case, there aren’t action-guiding standards. There are standards, but of a practically useless sort. I suggest, however, that there is at least one action-guiding standard. Consider the norm according to which it is impermissible to cause T-Pain. This is not a wildly complex disjunctive claim. It involves some disjuncts that are designed to eliminate limiting factors – specifically, it is wrong to deliberately cause pain unless it is deserved, or desired, or required as a means to a valuable end – but this is hardly unwieldy. It could easily play the role of an action-guiding standard. I therefore suggest that the Robust Realist can see necessary norms as playing an important role in ordinary life.

¹⁴ Sidgwick (1907) and Hare (1981). The example of two-level utilitarianism is also used by McKeever and Ridge though, it is worth emphasising again, it is just an example. The idea could be employed by a Kantian, a virtue ethicist, a care ethicist, and so on.

4.9 Conclusion

The claim that there are necessary moral norms plays a crucial role in the metaphysics of Robust Realism as I construe it, and in this chapter I have explored how these necessary norms can be understood. It's not plausible to see them as holding with conceptual necessity, as some suggest. More promising is the view that they are norms that have had their contingent limiting factors eliminated. Such moral norms are often, but perhaps not always fundamental. And this is how we can view the necessity in Robust Realism.

One might wonder how this form of necessity is to be categorised. It's not conceptual necessity, for it has nothing to do with the meanings of moral terms or the content of moral propositions. But can it be classed as a necessity of some other type? I'm not sure that identifying an easily labelled type of necessity would be helpful here. In a sense, the necessity in Robust Realism is nomic – it is to do with moral norms rather than natural laws, but moral norms may be laws of some sort. They are laws of morality, governing the behaviour of moral agents. In another sense, the necessity in Robust Realism is metaphysical – the moral norms would be just the same even with a different set of natural laws, and this is how metaphysical necessity is commonly understood. In yet another sense, the necessity in Robust Realism might be classed as a distinctive type of 'normative necessity' – it is a matter of there being norms that are *purely* normative, for they have had all contingent descriptive elements stripped away.¹⁵

You can use one of these labels if you like, but we don't really need to choose a specific classification. After all, what matters most is the idea, not the category or the label. And the idea is just that certain moral norms have modally maximal jurisdiction, because the range of worlds in which they have authority is not diminished by any limiting factor. Maybe other ways of interpreting necessary norms are available. I'd like to see further discussion, so that the ideas developed in this chapter can be compared against rival views. Still, we've made progress, and a neat account is now on the table. In the next chapter, I appeal again to necessary moral norms. This time, they play a role in explaining how the categorical authority of moral direction lends crucial support to the non-natural status of moral reality.

¹⁵ For the idea of a distinctive 'normative necessity,' see Fine (2002). I'm not sure my view really gels with what Fine has in mind, however. It is closer to Scanlon (2014).

5 The Authority of Morality

Of the commitments that constitute Robust Realism, there are two that I have yet to discuss. In this chapter I clarify these Realist Commitments, and show how one provides support for the other. The commitments are as follows:

Categorical Authority. The authority of moral direction is paradigmatically categorical.

Non-Naturalism. Moral properties are beyond the purview of the natural and social sciences.

My ultimate aim is to clarify how commitment to the former can motivate commitment to the latter. In other words, I intend to explain how Robust Realists can draw on the categoricity of moral direction in order to defend the claim that moral properties and norms force us beyond the natural world.

I start by examining Categorical Authority in depth. We can begin to understand this commitment by looking at the images that are used to evoke it. I discuss these in §5.1, to put a rough idea of categoricity on the table. But these images are considered obscure and inchoate by many, so I explore the notion of categoricity further in §5.2 and §5.3. My aim there is not to provide any sort of analysis of categoricity. Rather, it is to identify important facts about it, facts that will help us to get a clearer grip on it. I will explain how one important fact about categorical normativity is that, other things being equal, a moral agent is necessarily blameworthy – in a certain modally strong sense – if they culpably violate a categorical moral norm.

After clarifying Categorical Authority, I then assess how it matters for Non-Naturalism. There are various ways of distinguishing natural and non-natural properties in metaethics. One needn't assume that there is a uniquely correct way of drawing the distinction, but one has to be clear about one's chosen approach. I thus explain what I don't have in mind by the distinction in §5.4, and what I do have in mind by it in §5.5. In §5.6 and §5.7 I show how the naturalist position that stands the best chance of capturing categoricity cannot in fact do so. I discuss the accusation that categoricity is 'queer' in §5.8, and I conclude in §5.9 that the categorical status of moral direction helps to motivate a non-naturalist view of moral reality.

5.1 Images of Authority

Robust Realists, as I understand them, endorse the following view about the authority of moral direction:

Categorical Authority. The authority of moral direction is paradigmatically categorical.

Why claim that moral direction is *paradigmatically* categorical? After all, we could instead claim that it is always and essentially categorical, or that it is just occasionally categorical. However, both of those claims introduce certain problems. The problem with saying that moral direction is always and essentially categorical is that this is a very strong claim, and is thus very hard to defend.¹ The problem with saying that it is only occasionally categorical is that this might mean that the categorical parts of our moral framework could be eliminated without much loss to moral thought and discourse.

Saying that the authority of moral direction is ‘paradigmatically’ categorical provides a middle-ground between these extremes. The claim is that certain model cases – plausibly the fundamental moral norms – are categorically normative. It will be easier to defend this than the claim that all moral norms are always and essentially categorical. Moreover, making the paradigm claim will mean that categorical direction cannot be cut out of moral discourse without much loss, for it would entail the loss of some model cases. The middle-ground is thus a good way for Robust Realists to go.

What is categorical normativity? It is hard to rigorously answer this question. Philippa Foot (1972) famously suggested that claims about categoricity express a ‘fugitive thought,’ for the commitments involved in and entailed by such claims are notoriously difficult to identify and pin down. As an initial way in, it is therefore worth exploring certain images that philosophers have often used to evoke the notion of categoricity. Imagery can only get us so far, but it is a useful starting point. I’ll discuss two images. First, the image of *inescapable bondage*. It’s often said that, if a claim about what one ought to do is categorically authoritative, one is inescapably bound by it.² What this amounts to is up for debate, but it is a way into seeing the apparently unconditional character of categorical norms.

¹ Error Theorists make this claim, and it makes life hard for them.

² Foot (1972), Williams (1985) Garner (1990), Joyce (2001; 2006), and Cuneo (2014).

Consider that some directive claims have a sort of authority, even though it is conditional on some contingent desire or convention. Take this case: *if you want to impress your boss, you should offer to work late this evening*. This is a case of ‘hypothetical’ direction – the traditional contrast to categorical direction. Here the direction does have authority, but only for those who want to impress their boss. If no other relevant conditions are in play, then this direction will be escapable for anyone who lacks that desire. The thought is that categorical normative claims are not like this. Such claims are unconditionally normative and, as a result, inescapably binding.³

Jean Hampton puts this by saying that the authority of categorical direction is “independent of social or psychological contingencies” (1998: 96). As this suggests, categoricity has something to do with *necessity*. One might propose that a categorically authoritative moral claim will in some sense necessitate the performance of a certain act, or perhaps that failure to perform a certain act in some sense necessitates the suffering of some moral demerit. I expand on this shortly, but there seem to be two lessons to glean from the image of inescapable bondage. First, a plausible view of Categorical Authority will include the idea that moral authority involves a necessity claim. Second, however this necessity claim works, it should connect up to the unconditional character of moral direction, and should thereby give sense to an agent’s being ‘inescapably bound’ by categorical norms.

For another evocative image, consider the following from Richard Joyce:

One important feature of moral judgements ... is the degree of *practical clout* that they often purport to have in our deliberations and interactions: Calling an action “morally correct” or “virtuous” or “wrong” or “just” is (putatively) to draw attention to a deliberative consideration that cannot be legitimately ignored or evaded (2006: 57-58).

Several ideas appear to be at work in this image of *practical clout*. I have already addressed what is perhaps the main idea, for the claim that one cannot *legitimately* ignore a consideration with practical clout is similar to the claim that the authority of categorical direction is inescapably binding.

³ Foot (1972) argues that ‘systems of hypothetical imperatives’ provide the inescapability of categoricity whilst being ultimately conditional. I discuss this below.

But there is another idea suggested by practical clout, to do with deliberative weight. Ideally, our view of categoricity will illuminate why categorical claims tend to be experienced as stronger than other directive claims. In other words, the practical clout of a categorical norm consists partly in our giving it a lot of weight in our reasoning about what to do. I take it that the idea of categoricity is distinct from that of overridingness, and that the question of whether moral factors can ever be overridden by non-moral factors is not automatically settled by their categoricity. But it is clear that, in addition to being illegitimate, ignoring categorical moral direction is quite hard. It's certainly not impossible, for people can and do ignore it. But ordinary reasoners tend to give categorical direction a lot of weight in their practical deliberation. So, a lesson of the multi-faceted image of practical clout is that our view of Categorical Authority should ideally provide some insight into the weight that we give to categorical direction.

The pair of images sketched above expose the core of categoricity: in further developing this notion, we must shed light on the unconditional or 'inescapably binding' character of moral direction via a necessity claim, whilst also aiming to clarify the weight or 'clout' that such direction has in practical reasoning. With these considerations in mind, in the next two sections I look at Categorical Authority in depth.⁴ The hope is to develop a way of getting a decent grip on this commitment, and to thereby provide clarity to those who worry about the inchoateness of the images sketched above. I won't be offering a fully-fledged analysis of categoricity. I doubt that a deep analysis of it can be had, and I doubt that one is needed. Maybe a shallow ('intra-normative') analysis of categoricity is available, but its shallowness will mean that it has little interest. So, instead of offering an analysis, I aim to identify and clarify certain interesting facts about how categorical normativity relates to certain necessity claims. This will help us get a better grip the images outlined above, and will put us in a position to ask whether Categorical Authority supports a Non-Naturalist view of ethics.

⁴ Another image used in relation to categoricity is that of a 'demand without a demander.' For example, Garner (1990) uses this image in the course of claiming that categorically authoritative moral norms will be intolerably mysterious or 'queer.' I return to this claim in §5.8. Note as well that Anscombe (1958) uses a somewhat similar image of 'legislation without a legislator' in order to challenge the idea that moral obligation or duty make sense without God.

5.2 Categoricality and Practical Necessity

Like many others, I have said that categorical direction has something to do with necessity.⁵ In this section and the next I clarify the nature of the link between necessity and authority. One idea is that categorical direction necessitates the performance of a certain act. Like many others, however, I find this obscure. I will thus improve on it with something a little more careful.

Note first that, whilst our view of categoricality should illuminate the weight that categorically normative claims carry for the deliberating agent, we should distinguish the necessity in Categorical Authority from the necessity that some suggest is involved in moral motivation. In other words, we must not treat categorical normativity as a matter of an agent's being necessarily motivated to Φ by a moral norm that directs them to Φ . Moreover, we must not treat categoricality as a matter of an agent's being necessarily motivated to Φ by their judgement about such a norm. These are different issues, different sorts of necessity claim than that involved in categoricality. I don't know that anyone makes the mistake of conflating them – Mackie (1977) can seem to do so if interpreted uncharitably – but the necessity in Categorical Authority is not a motivational matter. So we cannot get a grip on categoricality by pinning it to this sort of motivational necessity.

More interestingly, one might try to get a grip on categoricality in terms of an agent's being such that a certain act seems unavoidable to them. The idea is that an agent's practical commitments – which go some way toward forming their character – are such that they regard an act as one that they must perform. The famous claim that has been attributed to Martin Luther – 'here I stand, I can do no other' – is supposed to illustrate the phenomenon. Although he probably didn't actually say it, the thought is that Luther's character and commitments were such that the option of recantation seemed (to him) practically unavailable. As a result, he could not escape his refusal to recant. This phenomenon is termed 'practical necessity' by Bernard Williams (1981a) and 'commitment authority' by Hampton (1998: 86). So, does this practical necessity help us to pin down categoricality?

⁵ This claim is made by those hostile to robust categoricality and by those sympathetic to it. For the hostile, see Foot (1972: 311), Williams (1985: 188), and Finlay (2014: Ch. 7). For the sympathetic, see Hampton (1998: Ch. 3).

Well, it might seem to help us to grasp both the inescapability of categoricity and the weight that it carries in our deliberations. After all, if you feel that you can do no other than Φ , then Φ -ing will seem inescapable to you. And the deliberative weight of categorical direction is explicable in terms of the strength of your practical commitments. So, one may think that we can get a grip on Categorical Authority by pinning it to practical necessity. However, this would be a mistake. Cases of practical necessity are interesting and important, but we mustn't tie categorical direction to them.

To see that there is a difference the two phenomena, consider that practical necessity provides a much weaker form of inescapability than can be found within categoricity. Even if one feels bound to Φ in light of a practical commitment, one could presumably escape the direction to Φ as a result of some change in that commitment. The change may be in the strength of the commitment, or in its relationship to one's various other commitments. Or the practical commitment might disappear altogether, perhaps to be replaced by another commitment. Categorical demands are not like this. They need not change in line with changes in one's practical commitments, and we should therefore avoid tying practical necessity to categoricity. The inescapability involved in practical necessity is weaker than the unconditional force of categorical direction. To view categorical direction through the lens of practical necessity is, therefore, to settle for less.

You might object that the depth of the practical commitments that generate practical necessity ensures that any changes in them will entail fundamental changes in character. If that's true, then one can only escape the requirements of practical necessity by becoming a different person. And this is a pretty strong notion of inescapability, you may think, for to escape the demands of practical necessity you must entirely change who you are. If this is true then, in the closest possible world at which Luther failed to do as he actually did, he was a different person. But would Luther genuinely have been a different person if he had done otherwise? I don't know for sure. I never met him. But it's very far from obvious. After all, if Luther had done otherwise, it would make sense for him to feel bad about himself. Presumably he couldn't feel bad about himself unless he was still himself.⁶

⁶ One might try to gloss his feeling bad about himself as a feeling of regret at no longer being the same man, but that seems implausible and potentially *ad hoc*.

It's also worth noting that the Luther example has the potential to mislead, for the practical commitments that made him unable to recant possess a certain social and historical significance, and involve a level of nobility and self-sacrifice. This might give this case of practical necessity the appearance of an authority that it doesn't in fact possess. After all, it is possible to imagine cases in which the commitments are far less noble, and thus less intuitively linked to categoricity. Imagine a neo-Nazi who is practically committed to the extinction of the Jewish culture. Even if these commitments generate practical necessity for him, this can only be seen as a psychological or motivational matter, and not as a normative form of inescapability, for we cannot tie the neo-Nazi's practical commitments to anything categorically normative. Put another way, even if one's practical commitments create a psychological inescapability, this gives us no grip on categorically authoritative normativity, for we can still ask whether or not we are categorically required to either cultivate or eliminate any of our practical commitments, including those that may generate practical necessity.

In sum, although certain forms of inescapability and deliberative weight may appear in the phenomenon of practical necessity, they do not help us get a grip on categoricity. We must look elsewhere to clarify this idea.

5.3 Categoricity and Necessitation of Blameworthiness

If practical necessity doesn't help, then what necessity claim should we make in trying to get a better grip on Categorical Authority? I suggest that, rather than looking to interpret the idea that an act is necessitated by a categorical norm, a nice way to proceed is to focus on what is necessitated by an agent's *failure* to perform the act prescribed by that norm. For here there is an obvious and attractive answer: *blameworthiness*. If an agent is categorically required to Φ , and if they nevertheless fail to Φ and are also culpable for that failure, then they'll necessarily acquire the property of *being blameworthy*. This is an inescapable aspect of categorical direction: one cannot escape the blameworthiness that comes from failure to follow it.⁷ The acquisition of this property is not conditional on a convention or desire.

⁷ One also cannot escape the praiseworthiness that comes from Φ -ing, though only if one does so for the right reasons. I'll ignore this, however, to avoid orthogonal complexities involved in the idea of acting *for the right reasons*.

More carefully, I suggest that the following is an important truth about categorically normative direction:

Necessitation of Blameworthiness. For any norm N , if N categorically directs one to Φ in a set of circumstances C , then an agent's culpable failure to Φ in C would necessitate their being blameworthy.

We can then appeal to this claim to get a grip on Categorical Authority. A defender of this Realist Commitment will hold that the authority of *moral* direction to Φ is paradigmatically such that, necessarily, one is *morally* blameworthy if one culpably fails to Φ . One cannot escape the acquisition of this moral property. Note that Necessitation of Blameworthiness is not an analysis of categoricity. It just expresses an important and plausible fact about it. It may even be a platitude, but by focusing on this fact we'll get a clearer view of the necessity in Categorical Authority. And, crucially, it will help us to develop a way of testing naturalistic approaches to moral reality.

To illustrate how Necessitation of Blameworthiness applies in the moral case, consider Watkins. Like all moral agents, Watkins is subject to a categorical norm that prohibits causing T-Pain, that is, undeserved and undesired pain that is not required as a means to any valuable end. But, being a ne'er-do-well, Watkins tortures someone just for fun. Moreover, he is aware of the fact that he is subject to categorical moral norms that prohibit such activity. He knows that he is culpably violating a categorical moral norm. The idea is just that, having culpably violated the norm that prohibits causing T-Pain, Watkins is *necessarily morally blameworthy*. No desires or conventions get him out of this. It is something that he cannot escape.

This is certainly not to suggest that Watkins will ever be subjected to those forms of social censure that commonly go with judging someone to be blameworthy; he'll escape punishment and ostracism if no one finds out that he was the one who did the deed. The idea, though, is that he has become *worthy* of moral censure. This is something that he cannot escape, for it is necessitated by his violation of a categorical norm. It was a moral norm that he violated, so his blameworthiness is of a specifically moral sort.⁸

⁸ I have nothing interesting or new to say about the nature of moral blame. For some prominent and recent discussions, see Watson (1996), Hieronymi (2004), Scanlon (2008), Wolf (2011), Wallace (2011), and Fricker (2016).

Again, this might seem obvious. But it helps us to clarify some key points about categorical normativity. First off, it helps to distinguish two forms of categoricity. To see this, consider that we can use Necessitation of Blameworthiness to understand a weak form of categoricity that one finds in certain non-moral norms, like the norms involved in the game of tennis. It's true that a tennis player might not be *morally* blameworthy if they culpably violate a rule of the game, but it is plausible that they are nevertheless blameworthy in another sense.⁹ They merit a certain type of censure – this is why certain penalties are used to meet rule-violations in tennis.

Moreover, these penalties are merited as a matter of necessity: once you violate a rule of tennis, you necessarily merit whatever penalty has been specified as corresponding to that particular rule-violation. You can't escape being worthy of the relevant penalty, whatever your desires, and even if you in fact go uncensured because the umpire did not notice your rule-violation. So, do the rules of tennis have the same authority as categorical moral norms? Foot (1972) answers affirmatively. She thought that inescapability is found in 'systems of hypothetical imperatives.' Such systems include the sets of institutionalised rules that constitute games, social clubs, etiquette, and so on. You can ignore the rules of etiquette, for example, but you remain subject to them. These systems of norms don't go away just because you ignore them, even if you don't care about them at all.

Foot thought that morality should also be viewed as a system of hypothetical imperatives.¹⁰ It might be an especially complex and messy system, and we generally attach more importance to it than we do to most other systems, but on Foot's account it is still just a set of institutionalised rules of the same basic sort as those involved in tennis. So, if categoricity is read in terms of Necessitation of Blameworthiness, doesn't this play into the hands of a non-robust view of morality – like Foot's? Not quite. We need to make a distinction between two types of categoricity. One is modally 'fragile,' whilst the other is modally 'robust.' This distinction maps on to one I made in Chapter 4, between limited and unlimited moral norms.

⁹ You may say that it is immoral because it is cheating. But one can break a rule of tennis, and be culpable for doing so, without cheating. For to cheat one must violate a norm on purpose, and one can culpably violate a norm without this doing so purpose.

¹⁰ At least, she did in 1972. Foot (2001) offers a different story.

Recall that a limited moral norm is one that in some way rests on or presupposes contingencies – ‘limiting factors’ – and that, as a result, has authority at some worlds but not at others. An unlimited moral norm is one that involves no limiting factors, and that, as a result, has authority with modally maximal jurisdiction. A system of hypothetical imperatives is not as limited as some views of the normative, but it is still restricted by its ultimate contingency on the rules being instituted at a world. The modal jurisdiction of the direction provided by a system of hypothetical imperatives is thus limited. Within this jurisdiction, direction provided by the system in question may be inescapable. So we can call it categorical. But this is a modally fragile form of categoricity, for there are worlds in which it fails to apply – worlds in which the rules are not instituted. By contrast, modally robust categoricity is inescapable at every world.¹¹ For there to be a modally robust categoricity, we thus need necessary norms. And, as we saw in earlier chapters, Robust Realists accept such norms. So they can appeal to them in interpreting Categorical Authority as modally robust.

In itself, then, Necessitation of Blameworthiness does not force us to choose either a modally robust or fragile view of Categorical Authority. This choice will be determined by debate about the existence of necessary moral norms: if there are necessary or unlimited moral norms, then we can have a modally robust form of categoricity, but if all moral norms are in some way limited and contingent, then the most we can hope for is a modally fragile form of categoricity. I have already argued (in Chapter 3) that Robust Realists can reasonably maintain that there exist necessary norms. So, whilst rival metaethical theories can draw on Necessitation of Blameworthiness, a robustly realistic theory is able to do so in a modally robust way that fits with a far more uncompromising view of categoricity than Foot’s. Moreover, a modally robust view of morality’s inescapability is no more obscure than the modally fragile view. All the modally robust view alters is the range of worlds in which a norm applies. You might deny the existence of unlimited moral norms, but you can’t accuse the modally robust view of categoricity of an obscurity that the fragile views somehow avoid.

¹¹ Foot might get this robust inescapability by employing a rigidifying manoeuvre. Even if this is so, it doesn’t change my distinction. Unlike robust views, fragile views must rigidify to get modally general inescapability. Given that Robust Realists already accept unlimited norms, they needn’t be hostage to the viability of a rigidifying move. Foot’s categoricity is also different to stronger forms because the latter is ‘reason-giving’ – cf. Joyce (2001).

So far we've seen how Necessitation of Blameworthiness is a fact about categoricity that helps us to get a grip on its inescapability: part of what is inescapable in categorical direction is the blameworthiness that is necessitated by culpable violation of it. This necessity claim comes in a modally robust or a modally fragile form, depending on whether there are necessary moral norms. Robust Realists will clearly opt for the modally robust approach. Again, this is not a deep analysis of categoricity. But it should help us to get a grip on its inescapability, and it will help us when we come to discuss naturalistic views of ethics. For we can test such views by seeing whether they are able to capture the claim that there are moral norms with an authority such that they necessitate, in a modally robust sense, the moral blameworthiness of moral agents who culpably violate them.

One final point before moving on. In addition to getting a grip on the inescapability of categorical direction, we also want to get a grip on the weight that it has in our deliberation. Although one can fail to do what one is categorically required to do, categorical direction has a practical clout that weighs very heavily for us. If we operate with Necessitation of Blameworthiness, we can explain this in terms of the negative feelings that (for most of us) go hand in hand with our judging ourselves worthy of blame. At the mild end of the spectrum one may feel embarrassment, whilst at the more serious end of the spectrum one might experience guilt or shame. These are feelings that most find unpleasant or distressing. We want to avoid them. The psychological impact of judging oneself worthy of blame should not be underestimated. If failure to Φ necessitates your being blameworthy, then you will probably give a lot of weight to Φ -ing in your reasoning, for being blameworthy is probably something that you dislike.

I haven't have analysed categoricity, but I have provided a way of getting a grip on it that improves on the evocative but vague images that are often found in the literature.¹² Failure to follow categorical direction necessitates blameworthiness. Now, in the moral case, there is a question about the scope of this necessity. Given that I have already explained why Robust Realists can and should postulate necessary moral norms, at this point we can allow that a modally robust view of categoricity is what we want.

¹² I've no doubt that some will still find categoricity mysterious and obscure. I consider a version of this charge in §5.8, but I've introduced enough clarity to push forward.

I'll soon argue that we have to go beyond naturalism to capture this modally robust view of categorical moral normativity, but note that Categorical Authority – even if it is interpreted as modally robust – does not *build in* a view of the natural or non-natural status of moral norms. If we wish to use Categorical Authority to defend Non-Naturalism, we must identify the naturalist view with the best chance of capturing categoricity, and argue that it cannot adequately do so. I make some such arguments in §5.6 and §5.7, but first let's discuss the natural/non-natural distinction itself.

5.4 Non-Naturalism and Metaphysics

Question: what is the correct way to distinguish natural and non-natural properties? *Answer:* there isn't one. In saying this, I don't mean to suggest that there is no reason for making the distinction. My point is that, whilst ordinary language places *some* constraints on the extension of 'natural,' it is unlikely to be uniquely definable. Its ordinary or everyday usage has various aspects, so it would be a mistake to regard just one philosophical account of naturalism as legitimate. In sketching my preferred approach, I am thus not denying that there are other interesting distinctions in the vicinity.

There are two broad ways of distinguishing between natural and non-natural properties. One is metaphysical, the other is methodological. The former approach seeks to identify a metaphysical commonality that characterises natural properties. If there are non-natural properties, they are those that cannot be characterised via this commonality. By contrast, a methodological approach seeks a methodology that is characteristic of our inquiry into the natural world. If there are non-natural properties, they are those that are beyond this form of inquiry. In this section I clarify why I don't take a metaphysical approach. That is, I explain why I will not try to find a feature of natural properties that all such properties have in common.¹³

¹³ Note that some views of this feature are dialectically awkward. For instance, Plantinga views naturalism as the claim that there is "no all-powerful, all-knowing, perfectly good creator of the universe; furthermore, there are no beings much like him" (2010: 251). Depending on what 'beings much like God' means, this view may tie naturalism too closely to atheism. In other words, it may turn all atheists into moral naturalists. Many atheistic metaethicists see themselves as moral non-naturalists. Plantinga's approach turns them into naturalists, so it is dialectically awkward. We should thus seek to provide a different view of the natural/non-natural distinction.

One metaphysical proposal is that the natural entities are those that are causally networked. That is, those that have some causal impact on the world. The non-natural properties, if there are any, are those that have no causal impact. One might think that this is an appealing way to view the natural/non-natural distinction in metaethics, for most Robust Realists do in fact see moral properties as causally inert. So, maybe moral properties are non-natural precisely because they are causally inert. The problem is that it is not clear that Robust Realists *must* see moral reality as non-causal – this was no part of the definition given in §1.1 – and there are some who do not. Graham Oddie has a robust view of value, but thinks that “values can affect us, causally, and it is through their causal impact on us that we can have knowledge of value” (2005: 2). Of course, this is a controversial claim. But it would be unwise to say that Oddie is neither a Non-Naturalist nor a Robust Realist simply because he takes this view of value.

Anyway, there is an even bigger reason to avoid the causal form of metaphysical naturalism. As many have noted, it fails to respect one of the few restrictions that ordinary language places on our account of the natural/non-natural distinction. I take it that competent speakers deny that we can view angels, demons, ghosts, and gods as part of the natural world. The same goes for the magical spells cast by wizards and witches. But such entities, if they exist, have a causal impact. So angels and ghosts (and so on) all count as natural entities if we take this causal approach to metaphysical naturalism. This is an unacceptable result, so we will have to look elsewhere in distinguishing natural and non-natural aspects of reality.

A second way of developing metaphysical naturalism is to identify the natural world with the physical world. On this sort of physicalist view, natural properties are those properties that are (or are reducible to) the entities posited by physics. So, the non-natural properties, if there are any, are those that are not (or are not reducible to) the entities posited by physics. Note that this is not just a methodological approach, on which a natural/non-natural distinction is drawn via the *methods* of physics. It is metaphysical, for what unites the natural properties on this view is not anything about the nature of inquiry in physical science, but rather the fact that natural properties are (or are reducible to) the various entities that physics posits. This physicalist view looks attractive, for it seems to provide a very neat way of distinguishing the natural and the non-natural world.

However, it also has serious flaws. Consider the famous dilemma that Carl Hempel proposed for physicalism.¹⁴ Physicalists understand the natural properties as those properties that are (or are reducible to) the entities posited by fundamental physics. But we might wonder whether they are talking about current physics, or the physics at which we'll arrive at the end of inquiry. Both options lead to serious issues. The first option is to draw on the entities posited by current physics. But if we go this way, physicalism is probably false. After all, current physics is incomplete. It is not done with its positing of properties, and there is a decent chance that its current posits will be radically revised in the future. Of course, this is more likely in some cases than others, but the history of science suggests that it would be a mistake to rely too heavily on assumptions about how our current physics will compare to future physics. So it would be unwise to rely on current physics. This is the first horn of Hempel's dilemma. Our other option is to appeal to the entities that will be posited at the end of physical inquiry. But if we go this way, then physicalism is vacuous. It can offer us no real guidance, for no one can say what entities will be posited by physics at the end of its inquiry. This is the second horn of Hempel's dilemma.¹⁵

In short, drawing the natural/non-natural distinction in terms of a physical/non-physical distinction is going to be tricky at best. It may be that there is a way of getting a clear physical/non-physical distinction on the table, but there is at present no way of doing so that is uncontroversial and clear. It is therefore unlikely to be useful in our discussion of moral reality. Other approaches to developing a strictly metaphysical account of natural/non-natural properties are no doubt available, but the two that I take to be most promising are not particularly helpful ways to go in this context. In the next section, I'll thus consider a methodological approach instead.

¹⁴ Hempel (1969). For a more recent attempt to carefully refine Hempel's dilemma, see Stoljar (2010: Ch. 5).

¹⁵ One might respond by saying that it is rational to believe that our current physics is correct and mostly complete. Alternatively, one might seek a way of understanding the physical/non-physical distinction that doesn't appeal to either current or future physics. Perhaps we can appeal to common sense, or a model case of a physical theory. I doubt that any of these proposals will work, but it would take me too far afield to go into it here. What matters here is that the physicalist/non-physicalist approach to the natural/non-natural distinction will at best need serious development to have any chance of being workable.

5.5 Non-Naturalism and Methodology

A methodological approach to naturalism seeks to identify a methodology or form of inquiry that characterises our discovery of properties that are generally taken to be natural. This approach means that, if there are any non-natural properties, they are those that cannot be discovered via this methodology or form of inquiry. I take a methodological approach to the natural/non-natural distinction. Specifically, I see the difference between them in terms of certain differences between scientific and non-scientific inquiry.

So, a property is natural if it is within the purview of a natural or social science, like physics, chemistry, biology, neuroscience, psychology, anthropology, economics, and so forth. A property is then non-natural if it is not within the purview of a natural or social science, that is, if the means by which we inquire into the existence and character of the given property is non-scientific. In other words, the natural properties can be investigated by certain *a posteriori* forms of inquiry that characterise the methodology of the sciences, whereas the non-natural properties cannot be investigated in this way. This means that moral properties are natural if inquiry into their existence and status is an empirical matter involving the methodology of the sciences, and they are non-natural if inquiry into their existence and character is at least primarily a non-empirical matter that must be pursued by a methodology other than those involved in the sciences.¹⁶

There are a number of ways in which both moral naturalism and moral non-naturalism can be developed, of course, even within this sort of methodological approach. Before exploring moral naturalism in more depth, however, I will first expand on my preferred way of drawing the natural/non-natural distinction, for whilst it is a common way of setting things up in metaethics, it does raise certain issues.¹⁷ For instance, if we distinguish natural and non-natural properties via the natural or social sciences, we might want to see a way of assessing whether a discipline is scientific.

¹⁶ I admit that I'm short on details as to what exactly this will be. This is because my project is in moral metaphysics, and I set questions relating to moral epistemology aside for now. I will, however, say that insights from *modal* epistemology are going to be relevant in a moral epistemology. For some Non-Naturalist epistemologies, see Shafer-Landau (2003), Audi (2004), Huemer (2005), Tropman (2008; 2009), and essays in Hernandez (2011).

¹⁷ The methodological approach goes back at least to Moore (1903). For recent uses, see Brink (1989), Darwall et al. (1992), Shafer-Landau (2006), and Cuneo (2007b).

I don't have a definition, and I'm not sure that one is available. But we have reason to see a discipline as scientific if it gives a central role to repeatable empirical experiments as a way of gathering data that either corroborate or falsify unifying or explanatory hypotheses about specified parts of reality. That is, scientific inquiry starts with a specific and often novel hypothesis, one that is in some way linked to a theoretical system of (in normal cases) well-supported core and auxiliary hypotheses. The specific and novel hypothesis aims to unify a set of disparate hypotheses within that theory, or to explain certain recently discovered phenomena that pertain to the relevant portion of reality. A scientist then performs repeatable empirical experiments that either corroborate or falsify their hypothesis.

This is certainly plausible as a rough gloss of what happens in the natural sciences, where the researcher is likely to be involved in the design of the experiment. But we can also suggest that it happens in the various social sciences, like economics, where the adoption of a certain economic theory in one set of societies and the adoption of another in a different set of societies may play the same sort of experimental role. Of course, this gloss ignores many important details about the natural and especially the social sciences. However, it will do here. I take it that any plausible view of science will put such methodological factors front and centre.

Another question: why draw on science and not empirical inquiry in general? David Copp opts to take that broader approach, on the basis that "we can have empirical knowledge that is not scientific, including knowledge of street names, dollar bills, aches and pains, and popular foods" (2003: 185). First off, note that the science-focused approach as I understand it does not *exclude* our being able to get empirical knowledge by non-scientific means. I said that natural properties can be known by the empirical methods of the sciences, not that they can only be known in this way. And it is clear that street names and such could be investigated by scientific methods – they are *within the purview* of the sciences – even if in practice we don't need to use scientific methods to get knowledge of them. But we still require a reason to tie natural properties to scientific inquiry, rather than to empirical inquiry more generally. My worry about an empirical approach is the same as my worry about a causal version of metaphysical naturalism. Both violate the same rare restriction that our ordinary use of 'natural' places on a philosophical account of the natural world.

Recall that competent speakers will deny that angels and ghosts (etc.) can be captured within the natural world. But such entities, if they exist, can have a causal impact. Angels and ghosts will thus be knowable empirically, via their effects. An empirical approach will thus treat them as natural, even though ordinary discourse does not. However, despite their causal impact, angels and ghosts (etc.) are not amenable to scientific inquiry.

Hypotheses about angels or ghosts are rarely as novel and riskily specific as scientific hypotheses, and angelic and ghostly hypotheses are likely to be part of unfalsifiable and therefore unscientific theories.¹⁸ Any experiments that could feasibly expose the existence of angels and ghosts are unlikely to be repeatable, for such beings do not reveal their presence upon the whim of our research into them, and it is not within our limited power to make them appear on such occasions. So, if we tie the natural to the scientific, we thus avoid treating such entities as natural. One might object that we could at least in principle get a falsifiable and fully scientific theory of, for example, ghosts. However, even if we are able to do this without losing something that is essential to the concept of a ghost, my judgement is that bringing ghosts into the purview of the sciences would in fact be sufficient to naturalise them. For this reason, a methodological approach does better to tie the natural to scientific rather than empirical inquiry.

Before moving on, let's consider how what I've said relates to the descriptive properties. Simplifying just a little, properties can be:

- (a) descriptive and natural;
- (b) descriptive and non-natural;
- (c) non-descriptive and natural;
- (d) non-descriptive and non-natural.

The first two categories are straightforward. Into (a) goes a property like *being pleasurable*, and into (b) goes a property like *being loved by the gods*. If you identify a moral property like goodness with the former, you are a reductive naturalist. If you identify it with the latter, you are a reductive supernaturalist. I discussed and rejected these reductive views in Chapter 3.

¹⁸ Popper (1959) was wrong to treat empirical falsifiability as *the* demarcation criterion, but it is still part of being scientific. This is consistent with the lesson of the Duhem-Quine thesis: we needn't automatically reject a theory because of a falsifying result.

The third category is more controversial, but it captures one way of understanding a naturalist view that I will discuss in detail in the next two sections. On this view, goodness cannot be picked out by any purely descriptive predicate, but it is still within the purview of the sciences. It therefore goes in (c). Robust Realists put goodness in (d), however, for they see it as both distinct from the descriptive and beyond the sciences. I have already argued – in Chapter 3 – that moral properties do not go in (a) or (b). I will shortly argue that they probably go in (d) rather than (c). But what matters for now is that, in denying that the moral properties are natural, Robust Realists are committed to the following:

Non-Naturalism. Moral properties are beyond the purview of the natural and social sciences.

Now we are in a position to ask why Robust Realists should treat moral properties or norms as non-natural, that is, as being beyond the purview of the sciences. As we'll see in the next two sections, an appealing answer is that the most plausible conception of moral naturalism remaining to us is unable to accommodate some important facts about categorical moral direction.

5.6 Beyond Naturalism: Part I

Appealing to categoricity as a way of motivating Non-Naturalism is a not new idea. Many have held that naturalistic theories are unable to capture the authority of morality. However, this suggestion is sometimes made rather quickly, and without much serious argument as to how a naturalist might try to capture a strong form of categorical normativity. Perhaps this is because some Non-Naturalists are unclear in their view of categorical normativity, relying on evocative but vague images. I, however, provided a specific way of getting an intuitive grip on categoricity: Necessitation of Blameworthiness is a crucial fact about categorical normativity, and it can help us to grasp the notion of inescapably binding authority. But it doesn't straightforwardly imply anything for the natural or non-natural status of ethics.

This is true even when we see categoricity as modally robust. And at this point in the dialectic, we are entitled to view it in such terms. So, our question is whether a naturalist view can capture moral norms whose violation by a culpable moral agent would necessitate that agent's being morally blameworthy in any possible world in which the agent so acted. I will show that the most promising form of moral naturalism cannot do this.

What form of moral naturalism is this? Well, not a Descriptivist form of the sort that many defend.¹⁹ In a Descriptivist naturalism, moral properties are identical to natural properties that can be picked out by purely descriptive predicates (like *being pleasurable* or *being desired*) rather than non-natural properties that can also be picked out by such predicates (like *being loved by the gods*). I will not pay attention to this Descriptivist naturalism here, for I rejected it in Chapter 3. After all, if we have to reject Descriptivism, then we'll have to reject both naturalist and non-naturalist forms of it. But there is one form of naturalism that may have the resources to capture Categorical Authority in modally robust terms.

The view I have in mind is a naturalistic style of moral realism.²⁰ Its core metaphysical claim goes as follows:

Synthetic Naturalism. There is a synthetic identity between moral properties and certain natural properties.

To say that there is a *synthetic* identity between the moral and the natural is to say that this identity is not implied by the conceptual content of our moral terms. Just as the identity between water and H₂O is not implied by the conceptual content of 'water' and 'H₂O,' the identity between moral property *M* and natural property *N* is not implied by the conceptual content of 'M' and 'N.' Because of this, these are identities that can only be discovered *a posteriori*. The most prominent attempts to defend this view can be found in the work of Richard Boyd, Nicholas Sturgeon, and David Brink.²¹

¹⁹ Smith (1994), Jackson and Pettit (1995), Jackson (1998), and Brown (2011).

²⁰ The view that I focus on below is the one that I think has the best chance of capturing Necessitation of Blameworthiness in modally robust terms. I don't deny that other forms of naturalism can capture Necessitation of Blameworthiness in modally fragile terms. For instance, Finlay (2008; 2014) develops a pragmatic approach to categorical language that may succeed on that more modest score. Moreover, it may well be that neo-Aristotelian views – like those of Nussbaum (1992), Hursthouse (1999), Foot (2001), and Thomson (2001; 2008) – can also give a modally fragile categoricity. A fuller discussion would explain why these styles of naturalism are unlikely to succeed in modally robust terms, but given limitations of space I'll focus on the naturalist view that forms the most direct competition to Robust Realism when it comes to robust categorical normativity.

²¹ Brink (1984; 1989; 2001), Boyd (1988; 2003a; 2003b), and Sturgeon (1988; 2006). The view is often called 'Cornell Realism,' but for mostly boring reasons I prefer to just call it 'Synthetic Naturalism.'

If we accept that M and N are related by synthetic identity, such that the terms used to pick out those properties are not interchangeable, then we can categorise moral properties as natural but non-descriptive.²² Although many treat the natural as a part of the descriptive, the Synthetic Naturalist claims that moral predicates and natural predicates diverge in their conceptual content. If this is true, moral properties cannot be picked out by any purely descriptive predicate. In the sense discussed in Chapter 3 – on which we treat the purely descriptive properties as those picked out by purely descriptive predicates – the moral properties are thus not appropriately understood as descriptive properties. Nevertheless, moral properties are natural properties because they are within the purview of the sciences. By this use of terms, moral reality is thus understood by the Synthetic Naturalist as being non-descriptive but natural. So, even though defenders of this view make an identity claim, it is not the sort of identity claim that Descriptivists make. Synthetic Naturalism is not a Descriptivist view.

Synthetic Naturalism was originally motivated by a certain moral semantics, an early version of which comes from Richard Boyd. Boyd saw moral terms as having their referents fixed in the following way:

Roughly, and for nondegenerate cases, a term t refers to a kind (property, relation, etc.) k just in case there exist causal mechanisms whose tendency is to bring about, over time, that what is predicated of the term t will be approximately true of k ... Such mechanisms will typically include the existence of procedures which are approximately accurate for recognising members of instances of k ... and which relevantly govern the use of t , the social transmission of certain relevantly approximately true beliefs regarding k , formulated as claims about t ... etc. When relations of this sort obtain, we may think of the properties of k as regulating the use of t (via such causal relations), and we may think of what is said using t providing us with socially coordinated epistemic access to k (1988: 195).

More crudely, the referent of a moral predicate is fixed by the natural properties that causally regulate its use. Some natural property causally regulates our use of ‘... is blameworthy,’ and this fixes the referent of this term.

²² They thus belong to (c) in the categories discussed in §5.5.

I don't know exactly what the Synthetic Naturalist should suggest as the referent of '... is blameworthy,' but what matters for now is simply whether they can capture the idea that culpable norm-violators acquire this property as a matter of unlimited necessity. That is, we want to know whether the Synthetic Naturalist's approach to morality is able to capture a modally robust view of Categorical Authority. Initially, it seems that it cannot. To see why this is so, consider Terry Horgan and Mark Timmons' classic 'Moral Twin Earth' objection to Synthetic Naturalism.²³ At the core of this objection is the thought that the use of moral terms can be causally regulated by different natural properties on different worlds. Consider two worlds – Earth and Moral Twin Earth – that are indistinguishable, save for the fact that their use of a moral term – 'blameworthy,' say – is causally regulated by different natural properties. It has the same sort of practical role on both worlds – both Earthlings and Moral Twin Earthlings use it to justify inflicting certain forms of censure on norm-violators – but for the Earthlings it is causally regulated by consequentialist factors, and for the Moral Twin Earthlings it is causally regulated by deontological factors.

The issue for Synthetic Naturalists is that it seems as though they must treat 'blameworthy' as meaning different things on Earth and Moral Twin Earth, even though this is implausible. After all, inhabitants of these worlds seem to *disagree*, but this would be impossible if they just mean different things by 'blameworthy.' Now, this objection can be developed in various ways, but what matters is the contingency that it reveals. Given that 'blameworthy' has distinct causal histories (and therefore distinct referents) in different possible worlds, a culpable norm-violator will not be necessarily blameworthy – in a modally robust sense – for their act. There will be worlds in which the use of 'blameworthy' is such that it would not apply to the norm-violator, even when it plays the same sort of practical role. This is a sort of interplanetary relativism, which may well slide into intraplanetary relativism.²⁴ So, if Synthetic Naturalists draw on Boyd's semantics, they cannot capture Categorical Authority in a modally robust form. At this point in the dialectic, Robust Realists have the upper hand.

²³ Horgan and Timmons (1991; 1992a; 1992b). For a precursor see Hare (1952: 148-149).

²⁴ These terms are from Brink (2001). Harman (2015) embraces this result as clarifying the nature of relativism, though most Synthetic Naturalists would want to avoid it.

5.7 Beyond Naturalism: Part II

A potential way forward for Synthetic Naturalists is to introduce a more subtle moral semantics, one that modally rigidifies the view so that it can avoid collapsing into some form of relativism. David Brink provides the two most promising rigidification proposals. The first starts with the idea that Boyd's semantics goes wrong in its focus on properties that causally regulate our *actual* use of moral terms. We will be better off if we instead see moral terms as referring to properties that regulate not just the actual use of these terms, but also the "counterfactual or hypothetical usage – in particular, the way speakers would apply terms upon due reflection in imagined scenarios and thought experiments" (Brink 2001: 168). This counterfactual view means that Earthlings and Moral Twin Earthlings can have a dialogue about the referent of 'blameworthy,' and it is possible that they will ultimately converge on a single way of using this term. Indeed, reflection may lead competent speakers to one use in all worlds. Perhaps, by taking this line, Synthetic Naturalists can see categoricity as modally robust.

This counterfactual approach to regulation does appear to be an improvement on what Boyd was proposing, but it is still not compelling. It might be a reasonably healthy optimism to assume that there will be convergence among Earthlings and Moral Twin Earthlings – they do live on almost indistinguishable worlds, so they start with a lot of common ground.²⁵ But it seems unreasonably optimistic to assume that competent speakers in *any* world would converge on one use of 'blameworthy.' Even if it does have the same referent at every possible world, it would take a great deal of argument to show that we can get universal convergence on it. I don't know exactly how unlikely it is that we'll achieve this universal convergence, but given that the range of worlds in question will differ in many more respects than Earth and Moral Twin Earth, pessimism is not unreasonable.²⁶ This is not to deny that the *possibility* of convergence may provide interplanetary dialogue with a sort of regulative hope. But this makes Synthetic Naturalism hostage to the fortunes of that discussion, for convergence in our use of '... is blameworthy' might nevertheless fail to occur.

²⁵ Laurence et al. (1999) argue that it will be hard to isolate those regulating properties, for different uses of *M* require many natural differences. So they may not have *that* much common ground. Even so, they have more with each other than they do with many worlds.

²⁶ For discussion of non-convergence in moral matters, see McGrath (2008; 2010).

Brink's second (and preferred) proposal starts with the shared referential intentions "that would justify us in interpreting a community of inquirers as engaged in moral inquiry" (2001: 172). And he thinks that the relevant referential intentions pertain to interpersonal justification. According to Brink, a predicate counts as moral if those who use it do so with the "intention of picking out properties ... that play an important role in the interpersonal justification of people's characters, their actions, and their institutions" (2001: 174). And, perhaps most significantly of all, this shared referential intention provides the semantic content of our moral terms.

To clarify, the thought is that '... is morally blameworthy' refers to certain properties – whatever they in fact are – that play an important role within the correct substantive moral view – whatever it in fact is – of interpersonal justification.²⁷ This allows Synthetic Naturalists to handle Moral Twin Earth by saying that Earthlings and Moral Twin Earthlings do disagree. They disagree about which natural properties make people and their actions and institutions interpersonally justifiable. Earthlings hold consequentialist views on this. Moral Twin Earthlings hold deontological views on it. Inhabitants of other possible worlds may hold very different views on it, but the meaning and referent of 'blameworthy' will not differ in those worlds as long as they share referential intentions with respect to it.

The contingency that was a problem before seems to be removed by a focus on shared referential intentions rather than causal regulation. True, inhabitants of other worlds may differ in their referential intentions with respect to predicates like '... is morally blameworthy.' But the point is that such speakers are not part of genuinely *moral* inquiry unless their referential intentions do relate to interpersonal justification, for genuine moral inquiry and discussion has interpersonal justification as one of its core constitutive aims. Can the Synthetic Naturalist thus deliver a modally robust approach to categoricity? I remain sceptical. A lot hinges on how exactly we cash out the notion of interpersonal justification. Brink doesn't go into much detail on this, but he does have some things to say. However, as we'll see, what he says causes problems for the Synthetic Naturalist's ability to interpret Necessitation of Blameworthiness in modally robust terms.

²⁷ Note that van Roojen (2006) takes a similar approach to Brink. Both rely on there being true substantive normative claims in developing their semantics.

In developing the idea of interpersonal justification, Brink draws on some remarks from Hume. The most relevant tells us that for a person to engage in moral talk he “must choose a point of view, common to him with others; he must move some universal principle of the human frame, and touch a string to which all mankind have an accord and sympathy” (Hume *Morals*: §9.1). This is nice, but it does raise questions. What is the status of Hume’s *must*? What is its modal jurisdiction? Brink’s spin on the Humean approach to morality fails to clear this up, for he only says that “what is distinctive of the moral point of view is that we assess people, and actions, and institutions according to standards that others can and should accept” (2001: 174). Brink’s remarks just raise the same questions as Hume’s: what exactly is the nature and modal jurisdiction of Brink’s *should*.

We can probably assume that Hume’s *must* and Brink’s *should* are practically normative.²⁸ This leaves us with the question of their modal jurisdiction. As we know, two types of answer are available. One is that Hume’s *must* and Brink’s *should* are limited in their modal jurisdiction. The other is they are unlimited in their modal jurisdiction. Let’s start with the former option. A simple version of this idea involves reading Hume and Brink as making conditional claims: *if you want to participate in moral discourse, you must/should intend (by your use of moral terms) to refer to properties that play a role in interpersonal justification*. If we go this way, however, we cannot view Categorical Authority as modally robust. After all, if categoricity is modally robust, then it is authoritative for all. Not just for those who want to participate in moral discourse. Of course, there are more sophisticated views than the above conditional claim, but the worry generalises. Any modally limited account will only support a fragile view of categoricity, for it won’t speak to those who refuse to engage in moral inquiry despite having a capacity for moral agency. Hume dealt with this by arguing that justice is “in the true interest of each individual.”²⁹ But this seems optimistic, if we have to include the moral agents of all possible worlds.

²⁸ This is how Rubin (2015) reads them. I came to Rubin’s article after developing the ideas below, but our objections to Brink operate with a similar strategy of questioning his theory’s ability to deliver a suitable underpinning for interpersonal justification.

²⁹ This is how Hume (*Morals* §9.16) replies to the ‘sensible Knave’ problem. Of course, he wasn’t seeking modal generality – the circumstances of justice obtain only in some worlds. Hume would not have worried about getting the robust authority that we’re after.

Many modern Humeans will be happy to take a different sort of line, and embrace the result that moral direction is escapable for some agents.³⁰ But that would be to abandon the goal of naturalising a modally robust reading of Categorical Authority. Perhaps the Synthetic Naturalist can be fine with that. I don't know if Brink himself would be, but (more pertinently) it is not what we want at present. At present we are asking if naturalists can provide a modally robust view of Categorical Authority, and they clearly cannot do so if they pursue this sort of modern Humean line. So let's assume that Synthetic Naturalists will deny that moral norms are limited in their modal jurisdiction, and that they will go the other way instead.

The other way for them to go is to interpret Hume's *must* and Brink's *should* as being unlimited in their modal jurisdiction. In other words, the authority of their practically normative claims – that moral agents *must* touch a string to which all humanity has an accord and sympathy, or that they *should* appeal to standards acceptable to others – extends to every world. It is inescapable for anyone with the capacity for moral agency – they must engage in the moral project of interpersonal justification. This would provide our modally robust account of categoricity. The demands made within moral discourse would not be contingent on any constitutive aims of that discourse. Such demands would be categorically binding, for even if those with the capacities constitutive of moral agency were to fail to engage in moral discourse, they would necessarily be at fault for doing so.

However, this also raises some difficult questions. If one's goal is to provide a naturalistic explanation of the modally robust categoricity of moral direction, one cannot just appeal to the modally robust categoricity of non-moral direction without raising questions about how that modally robust non-moral categoricity is itself to be understood. In other words, Synthetic Naturalists must now provide a general account of categorical normativity. But this creates problems, for there are two paths available to Synthetic Naturalists and neither are promising. The first path involves providing a view of categorical normativity in Synthetic Naturalist terms. The second path involves appealing to some alternative meta-normative theory.

³⁰ Sophisticated versions of this modern Humean line can be found in Lenman (2010) and Street (2009; 2012).

Start with the first path. The idea is to find referential intentions constitutive of categorical normative discourse, and to then treat them as securing the referent of categorical normative terms. Such terms refer to whichever natural properties play the role fixed by the shared referential intentions of the participants in that discourse. I am not sure of what the relevant referential intentions will be, but I am sure that any workable proposal faces serious issues. This is because, to be workable, a proposal must have normative content built into it. Brink's appeal to interpersonal justification seems neat for the moral case because it gives moral norms a deeper normative underpinning. To work for categorical normativity in general, we'd therefore need some deeper normative underpinning for the normative underpinning offered by Brink. And this is where problems emerge.

After all, once that deeper normative underpinning is provided, the same questions that arose for the shallower normative underpinning will just re-emerge. Specifically, we have to ask for information about the nature and modal jurisdiction of that deeper normative underpinning. If it is modally limited, we lose the modal robustness that we're after. If it is modally unlimited, we must provide another normative underpinning for it. In that case, regress threatens – we will have to keep coming up with normative underpinnings. Either way, the first path is not a promising one for the Synthetic Naturalist. It cannot be used to vindicate a modally robust interpretation of Categorical Authority. So, what about the second path? The second path involves looking to some other meta-normative view of categoricity as providing the normative underpinning for ethics. Moreover, the view that we choose must support a modally robust view of categoricity. There aren't many options here, but one involves building Synthetic Naturalism about moral norms on a more generally non-natural base.³¹

³¹ Perhaps another is the Kantian constructivism defended by Korsgaard (1996; 2009) and others. I'm not sure if this constructivist view is, in the sense relevant here, a form of naturalism or non-naturalism. If it is a form of naturalism, perhaps Synthetic Naturalists can build on that meta-normative view to avoid having to accept non-natural entities. However, we should note that to build Synthetic Naturalism on any rival meta-normative theory would be to accept a problematic disunity. So, even if Kantian constructivism is a naturalist meta-normative view, and even if Synthetic Naturalists can coherently adopt it as a normative underpinning for their view of morality, the second path is still not an attractive way for them to go, for it would force them into a strangely disunified theory.

That is, Synthetic Naturalism at the moral level could be sustained by the idea that Hume's *must* and Brink's *should* are non-moral but also practically normative terms, both of which are unlimited in their modal jurisdiction, where this commits us to the existence of something beyond the purview of the sciences. But once one says this, it is hard to see why one would not see categorical moral norms as non-natural too. First, this provides a more unified account of the normative domain. Second, if one aim of the Synthetic Naturalist is to avoid putatively mysterious entities that exist beyond the purview of scientific inquiry, then the second path is unpromising for them, for such entities are likely to be its ultimate end point.

In sum, attempts to capture a modally robust form of categoricity in Synthetic Naturalist terms are likely to fail. Some ways of developing this view lead only to a modally fragile view of categoricity. This may be acceptable to some, but it is not what we want at present. Other ways of developing it are modally robust, but rely on a deeper non-naturalism. In short, to capture Categorical Authority in modally robust terms, we must deny moral naturalism and instead accept a Non-Naturalist view of moral reality.

5.8 Arguing from Queerness

If what I've said so far is correct, then a modally robust interpretation of categorical moral normativity will force us beyond the natural and social sciences. But one can accept this without agreeing that such norms exist. For example, Error Theorists can agree that moral discourse is committed to categorical moral norms, but they still think that there simply are no such norms.³² One of the most common accusations from Error Theorists is that categorical moral norms, at least when viewed as modally robust, would be metaphysically *queer* entities. According to Error Theorists, we should therefore avoid positing them, even if they are an essential part of moral discourse. The upshot will be that moral discourse is systematically infected with an error, because it is committed to the existence of queer entities. In this section, however, I will explain that it is not dialectically helpful to argue from queerness, for Robust Realists do not find morality queer.

³² There are routes to Error Theory that sidestep categoricity and focus instead on moral supervenience, moral knowledge, or moral motivation. It is ultimately more promising to look to categoricity, however. Cf. Garner (1990), Joyce (2001), and Olson (2014).

A lot hangs on how we are to interpret the idea of queerness itself. J.L Mackie's original claim was that an entity's queerness is somehow a matter of its being "utterly different from anything else in the universe" (1977: 38). This can make it seem as if the queerness charge is basically a charge of uniqueness. Understood in this way, however, it isn't going to trouble any Robust Realists who follow the recommendation offered at above, and build into their theory the view that at least some moral norms are *fundamental* aspects of reality. After all, there is an obvious sense in which all fundamental entities, including those of physics, are unique. Nothing else does exactly what a quark does, and nothing else is quite like a lepton. In short, even if nothing else in the universe works quite like a categorical moral norm, in itself this is not enough to undercut belief in them.

However, queerness is not best understood as uniqueness *per se*. Rather, it is the way in which moral norms are unique that makes them queer. According to Error Theorists, it is categoricity itself that is peculiar. At this point, however, things get rather murky. In particular, we must ask Error Theorists just *who* it is that finds categorical normativity queer. It doesn't seem queer to me, or to most Robust Realists, so how exactly is a queerness charge meant to have bite against those of us who need to be swayed? Maybe an Error Theorist can reply by agreeing that categoricity won't be obscure to everyone at first glance. They just claim that a careful examination of the idea of a categorical norm reveals it to be intolerably inchoate, such that we must treat it as committing moral discourse to an error.

But that's not a compelling move now, for I have already shown that there is no inherent obscurity even in the idea of a modally robust categorical moral norm. The point of Necessitation of Blameworthiness is that it helps at least some of us to get an intuitive grip on the notion of categoricity. It may not help an Error Theorist, but it does give me a clear grasp on the notion of categoricity.³³ Even if the necessity here is modally robust, there is nothing obscure about it. It's just like the modally fragile view, except applied to a wider range of worlds. Both the Robust Realist and the Error Theorist should thus agree that categoricity, understood in this way, is not inherently obscure. So, how can the Error Theorist sway those who still don't see the queerness that they see in categorical moral norms?

³³ This may be due to differences in temperament, an idea that I explore in Chapter 6.

At this point, an Error Theorist might fall back on their own vague but evocative imagery. After all, certain Error Theorists have made their queerness arguments by appealing to images designed to evoke concern about the notion of a categorical norm. One of the best-known is that of a ‘demand without a demander,’ which some have used in an effort to sway those who believe in categorical moral norms. The thought is that, whilst other types of demand (like legal and social norms) derive their authority from identifiable demanders (such as legislators and communities), this is not true of categorical norms. Such norms do not derive their authority from anywhere else, and are thus a lot queerer than those that do. Richard Garner, for example, says that whilst we “know what it is for our friends, our job, and our projects to make demands on us, we do not know what it is for *reality* to do so” (1990: 143). So, perhaps critics of categoricity can use this sort of imagery to sway the Robust Realist against categorical norms.

In response to this strategy, I suggest that it is both unhelpful and uncharitable to say that, by the Robust Realist’s lights, it is *reality* that makes demands on us. This makes it seem as though the Robust Realist is anthropomorphising reality, as if they are treating it as an agent with the capacity to give instructions, prescriptions, recommendations, and so on. This would of course be a bad move, and perhaps ‘queer’ is an apt term for the manner in which would be a bad move. However, when the Robust Realist says that there exists a categorically authoritative moral reality, they are not suggesting that reality itself somehow makes demands of us. They are just saying that categorical moral norms exist as part of reality. Of course, Robust Realists may be wrong in their claim that robust moral properties and norms do in fact exist. But it is a mistake to interpret their view in the way Garner does, and it is more generally a mistake for the Error Theorist to develop the queerness charge in this way. Imagery may have legitimate uses in the back and forth between Robust Realists and Error Theorists, but not if it ends up misleading us by caricaturing either view.³⁴

³⁴ In fairness to those who present an uncharitable interpretation of the Robust Realist’s commitment to categoricity, it is not often that the idea of a categorical moral norm gets developed in the rigorous sort of way discussed earlier in this chapter. Even so, the image of a demand without a demander is still unhelpful, and having put in place an account of categoricity – Necessitation of Blameworthiness – there is no serious pressure on me to provide an alternative. I have already proposed a rigorous interpretation of categoricity that involves no assumption of reality making demands.

However, my strong suspicion is that, in the end, arguments about the existence of robustly categorical norms are not going to be settled by intuitions pertaining to their queerness, or lack thereof. Error Theorists see something queer in robust categoricity, and are therefore unlikely to be persuaded that it exists. Robust Realists don't see something queer in robust categoricity, at least not when it is seen as part of the fundamental structure of reality. They are therefore unlikely to be swayed by a charge of queerness. Jonas Olson is right to say that the queerness debate puts Robust Realists and Error Theorists "in a stalemate, staring incredulously at each other" (2014: 136). I would go further, and say that many debates between realists and anti-realists very often result in stalemate, and that this general stalemate needs some diagnosis. I develop this in detail in the next chapter, looking at how we might overcome the stalemate and move forward.

For now, however, we can just note that categoricity cannot easily be dismissed on the grounds that it is queer. If the charge is just that it is a unique aspect of the universe, then this just means that the fundamental categorical moral norms are on a metaphysical par with quarks and other fundamental physical entities. If the charge is just that there is something odd about categoricity itself, then it won't sway those who don't see this oddness. In sum, Robust Realists are unlikely to be swayed by a queerness charge.

5.9 Conclusion

Insofar as we are entitled to prefer a modally robust view of categorical moral norms to a modally fragile one, Categorical Authority helps us both to understand and motivate Non-Naturalism. Of course, this won't worry those who reject modally robust views of categoricity. I haven't shown that such views are definitely false. That is, I haven't argued that moral discourse is committed to there being modally robust categorical norms. This is primarily a semantic rather than a metaphysical question, and as my present focus is metaphysics I'll set that issue aside for another time. Nevertheless, in earlier chapters we saw that one can reasonably believe in modally unlimited moral norms. If we are entitled to believe in such norms, then we are entitled to take a modally robust view of categoricity. And in this chapter we've seen that this modally robust view will force us to go beyond of the natural and social sciences in understanding moral reality.

6 Arguing about Moral Reality

In all that I have said so far, my goal has been to interpret and motivate a metaphysics for Robust Realism. I've developed a cluster of metaphysical claims that should be endorsed by Robust Realists, and that have enough support to ensure that even critics must take them seriously. However, what I've not yet asked is whether it is Robust Realists or their critics who are actually right. I haven't given an argument for the claim that there are any fundamental moral norms that are at once irreducible, non-natural, mind-independent, and categorically authoritative. I *have* said that belief in robustly necessary moral norms is at least reasonable – there is some evidence for it, and the existence of robust moral norms cannot be quickly dismissed as unacceptably queer – but I *haven't* shown that they do in fact exist.

Nor do I intend to do so now, and not because I'm inclined against Robust Realism. Rather, because I judge that arguments for it are unlikely to convince those who need convincing. In this chapter I show that debate about moral reality faces major and perhaps insurmountable obstacles. I begin in §6.1 by outlining a common way of arguing about moral reality. In §6.2 I show how it generates an intractable stalemate. I diagnose the intractability of this stalemate in §6.3 and §6.4, showing how differences in temperament play a major role in its entrenchment. We thus need ways to avoid or overcome the influence of temperament. There isn't space to look at every way of trying to do this, so I focus on just one argumentative strategy that might seem to move the debate forward at the expense of Robust Realism. I outline this evolutionary 'debunking' argument in §6.5. In §6.6 and §6.7 I show that it fails to avoid or overcome temperament. I consider how the future of metaethical inquiry can best proceed in §6.8, concluding in §6.9 that it may have to operate with surprisingly humble ambitions.

6.1 The Substantive Strategy

There is a common way of arguing about moral reality that generates an intractable stalemate. I explain this way of arguing in detail below, but it basically involves assessing metaethical views against their implications at the substantive moral level. So let's call it 'The Substantive Strategy.' I sketched some examples of this style of argument in earlier chapters, and in this section I look at further examples to carefully expose its central features.

I'll look mainly at how The Substantive Strategy plays a role in the debate about Mind-Independence, for this is the aspect of Robust Realism to which I have paid the least attention. Recall that it says the following:

Mind-Independence. Moral properties are constitutively independent of any agent's or set of agents' actual or hypothetical attitudes, beliefs, and conventions.

To illustrate how The Substantive Strategy is used to defend this claim, let's start with the examples that we witnessed in earlier chapters. Recall, from §2.8, that Relaxed Realists try to defend the Realist Commitments by appeal to substantive moral intuition. Matthew Kramer maintains that mind-dependent views of moral reality – according to which moral truths depend on the attitudes of agents – are untenable because they imply intolerable views at the substantive moral level. Specifically, they imply that there are circumstances in which it is permissible to torture babies, for this is something that an agent might coherently desire. Kramer finds this repugnant, and says that it undercuts mind-dependent theories. For Kramer, the debate about moral reality and Mind-Independence is to be decided by examination of how theories fare against substantive moral views.¹

Another example is from §4.7, in which I discussed a community whose conventional agreements permit the annual sacrifice of a teenager. Many will judge this community to be mistaken about what is right and wrong, even if their conventions are consistent, and even if they rest on no false non-moral beliefs. This substantive moral intuition is then used to argue that we must go beyond our attitudes (and so forth) to get to the correct moral norms. It is worth noting that this community is akin to an oft-discussed character in metaethical debate. Imagine someone whose desires (understood in the broad sense, as 'pro-attitudes') are internally consistent and whose non-moral beliefs are all true, but who has a desire for something that seems so odd or offensive that we consider them at fault.² Such figures will be familiar to those acquainted with metaethical debate. Sharon Street (2009) calls them 'Ideally Coherent Eccentrics,' or 'ICEs.'

¹ Kramer (2009). Cf. Dworkin (1996; 2011) and Nagel (1997).

² Clarifications. (1) Desires are internally consistent when the satisfaction of one does not preclude the satisfaction of any other. (2) It may be better to say that the agent has no false non-moral belief whose correction will alter their desires, not that their non-moral beliefs are true. This will depend on how much we should expect of agents.

This label is apt, for despite being ideally coherent in their desires and non-moral beliefs, these characters are eccentric in what they desire. So much so that they appear to make a substantive moral error, despite the fact that we can find no internal fault in the practical standpoint from which they deliberate and judge. ICEs often appear in debate about how our reasons relate to our attitudes. Here are some representative cases:

Adil is ideally coherent, but he prefers the destruction of the world to the scratching of his finger (Hume *Treatise*: §2.3.3).

Bert is ideally coherent, but he has no desire to avoid a future agony he is scheduled to suffer (Parfit 2011 v.1: §11).

Caligula is ideally coherent, but he aims solely to maximise the suffering of others (Gibbard 1999: 145).

Doris is ideally coherent, but she accepts norms prescribing starvation even if the result is a figure that cannot sustain life (Gibbard 1990: 171).

Edwina is ideally coherent, but her only pleasure is to count blades of grass (Rawls 1971: 432).

Francis is ideally coherent, but he is a Mafioso who values strength and honour above all else (Cohen 1996: 183).

These cases, and others much like them, are extremely common in ethics and metaethics.³ They are very often employed in arguments about moral reality.

To see how ICEs are employed, let's start with attempts to defend Mind-Independence, and let's focus on Caligula. Caligula's sole aim is to maximise the suffering of others. And, according to many defenders of Mind-Independence, it is intuitively obvious that this aim is immoral. Yet we cannot fault Caligula for it if he is ideally coherent and if moral reality is mind-dependent. We could fault him if his desires were inconsistent, or if some of his non-moral beliefs were false, but by stipulation they're not. So, if moral reasons derive from desires, and if Caligula is ideally coherent in desiring to maximise the suffering of others, we cannot derive a reason for him not to do so. So, to vindicate the substantive intuition that Caligula is morally mistaken, we must invoke some mind-independent reason or norm.

³ For many famous examples, see Parfit (1984). Note that some cases appear in debate about *moral* reasons, others in debate about reasons generally. Moreover, some see their impact as an ethical and not a metaethical matter. I see it as metaethical – see Chapter 2.

That, at least, is what many defenders of Mind-Independence who employ this version of The Substantive Strategy suggest. They think that intuitions about Caligula constitute evidence for Mind-Independence, for they reveal that mind-dependent theories are extensionally inadequate when it comes to capturing the substantive moral data. So, perhaps the Robust Realist can make use of such cases as part of an argument for their account of moral reality.⁴ How can mind-dependence theorists reply? A standard reply is to bite the bullet, whilst claiming that the bullet is quite edible.

That is, many mind-dependence theorists have denied that ICEs provide evidential support for Mind-Independence, on the basis that their examination of characters like Caligula yield intuitions that are instead friendlier to a mind-dependent view of moral reality. They say that, for someone as eccentric as Caligula to genuinely count as ideally coherent, he must have a bizarre psychology, or have been raised in a very strange environment – somewhere utterly unlike any typical human community. According to Street, for example, once we fully grasp the psychology and history that Caligula must have if he is to count as ideally coherent, he will seem more like an *alien* than a human being.⁵ And the same goes for other ICEs.

The mind-dependence theorist then claims that an appreciation of the strangeness or alien-like nature of ICEs ensures that it is no longer clear that such figures are mistaken or at fault. Indeed, many of them find it more intuitive that they are not mistaken. Street, for example, says that this is the case with Caligula. She suggests that, intuitively, he does have most reason to maximise suffering. And this is said to be evidence for a mind-dependent view of morality. After all, the intuition that Caligula has most reason to maximise suffering fits perfectly with the view that moral reasons derive from desires. So, the mind-independence theorist's reply involves a readiness to tolerate or embrace a certain substantive view of Caligula.

⁴ Robust Realists do sometimes use moral claims, despite my definition in §1.1, on which they are to primarily appeal to metaphysical factors. See, for example, Oddie (2005: 106), Huemer (2005: 50), Shafer-Landau (2004; 2009), and Enoch (2010a; 2011: Ch. 2).

⁵ Street (2009; 2012). For related claims, see Harman (1975: 7-8), Hare (1979), Williams (1981b), Lewis (1989), Shemmer (2007), Goldman (2009), and Lenman (2009; 2010). A different approach is to deny that ICEs are possible. Maybe *ideal* agents are logically or rationally committed to certain desires. See, for example, Smith (1994) and Korsgaard (1996). I find this implausible, but it won't matter for my argument, as I explain in n.8.

Note that the suggestion is *not* that cases involving ICEs are too far-fetched for our intuitions about them to be reliable.⁶ Instead, the claim is either that they are psychological outliers about whom it is acceptable to bite a bullet, or else that they in fact present no real bullet to bite, for a careful examination of them *supports* a mind-dependent view of ethics. This won't convince many defenders of Mind-Independence, I expect. For them, the intuition that Caligula has an immoral aim is likely to persist. I explore this in §6.2, but what matters for now is that cases involving ICEs illustrate how arguments about Mind-Independence – from its critics as well as its defenders – are often run in accordance with The Substantive Strategy. At the heart of this strategy is the thought is that we can assess metaethical theories by looking at their substantive moral implications. A good theory vindicates substantive moral intuitions. If a theory implies an unintuitive moral claim, this counts against it. We may disagree about what is unintuitive – more on this later – but the strategy is to defend metaethical conclusions by appeal to arguments with substantive moral premises.

Many metaethicists buy into The Substantive Strategy, and it does have attractions. Most metaethical theories imply the same answers for everyday cases, but when ideal coherence and eccentricity are stipulated, they divide. Attention to cases involving one or another ICE can thus seem crucial in metaethical debate. Street even says that such cases “are where the action is if we want to get clear on the relation between our attitudes, value, and the world” (2009: 279). Still, the notion of an ICE is just a vivid illustration of the broader style of argument that I call The Substantive Strategy. This strategy manifests in various ways and in various debates.⁷ Rather than seeing how else The Substantive Strategy can be developed, however, I'll move on to showing how it leads to a stalemate. The major problem with it is that we differ in moral intuition. We are prepared to accept different moral views, and defend different metaethical views as a result.

⁶ For discussion of this point, see Wood (2011) and Elster (2011).

⁷ For instance, some say that Mind-Independence is needed to capture the moral intuition that some acts are not just wrong but *evil* – see Russell (2006). ICEs need not feature here, but this is an example of The Substantive Strategy. A related case is the ‘isolation test’ developed by Moore (1903) and used recently by Wielenberg (2014). For a metaethical debate that features ICEs without using them to defend Mind-Independence, consider the amoralist in debates about moral motivation – see Brink (1989) and Smith (1994).

6.2 Stalemate

An implicit lesson of §6.1 is that it is possible for two competent inquirers to diverge in their substantive moral intuitions about an ICE, even if they are informed about the details of the case and even if they have given it careful thought. In this section I argue that this persistent difference in intuition (among seemingly competent inquirers) makes The Substantive Strategy a rather problematic way of arguing, for it leads to an intractable stalemate.

ICEs like Caligula are used by some to argue that moral reality is mind-independent. But people differ in their moral intuitions about this case. Some are happy (or at least willing) to accept that there is nothing immoral about Caligula's aim, others see it as intolerably repugnant. How likely is it that one side will shift in their intuitions? Not very, for those who argue over ICEs are intelligent and informed philosophers who have reflected carefully on the relevant cases. They just persistently disagree about them. Moreover, it would be uncharitable to suppose that one side is being dishonest or unreasonably stubborn, clinging to a false view from some intellectual vice, like pride. It's more likely that they are just arguing about difficult cases. In these difficult cases, reality does not present itself in a clear-cut way. If it did, reasonable and informed people who thought about them carefully would come to similar views. But it doesn't, so they don't.

ICEs are thus *controversial*. Not just because we disagree about them – after all, we can disagree over p even if it is obvious that p – but because conscientious and informed inquirers disagree about them even after thorough reflection on the same considerations (cf. McGrath 2008). Some might be very confident about their intuitions concerning ICEs, but the fact that other competent inquirers come to the opposing view should make even the most confident among us see them as controversial in the relevant sense. The crucial point to which I wish to draw attention relates to the dialectical effectiveness of philosophical arguments. Suppose that you are a philosopher with an argument for X , and that a key premise in your argument relies on an intuition about a case. The argument will only convince your audience of X if they share that intuition. If they don't, they won't see your argument as successfully supporting X . Moreover, if their intuitions are directly opposed to those required to make the argument go through, they are in fact more likely to see it as supporting the *denial* of X .

I suggest that this is what happens when Mind-Independence is assessed by appeal to ICEs. One side argues for Mind-Independence using the substantive intuition that, say, Caligula has no reason to maximise suffering. But this only persuades those who share that intuition, and see its capture as vital. It turns out that lots of competent inquirers don't feel this way. After careful thought, they are fine with saying that Caligula *does* have reason to maximise suffering. The result is stalemate.⁸ Defenders of mind-dependent and mind-independent theories of morality persistently fail to persuade each other when their arguments rely on intuitions about cases involving ICEs, for their intuitions about these cases persistently differ.⁹ Of course, it's not as if metaethicists are just impotently asserting their intuitions at each other. Sophisticated arguments are made on both sides of the divide. Still, when the arguments are reliant on controversial cases about which we persistently differ, they are dialectically ineffective. They fail to exert rational pressure on their primary targets, namely those not yet convinced of their conclusions. The Substantive Strategy is thus a problematic way of arguing about moral reality, for it leads to a stalemate. Given its position in current debate, the problem is not an inconsiderable one.

In §6.3 and §6.4 I diagnose this stalemate, explaining how it gets entrenched. Before that, however, I should clarify a couple of things to pre-empt some potential issues. First off, I'm not saying that there is no use for controversial cases. After all, it can be worth making an argument even if it doesn't speak to everyone. Such arguments might just preach to the choir, but some people haven't given it any thought, and need to be shown that they are members of the choir. Another use relates to the fact that we can often elucidate a theory's commitments by comparing cases with one another (cf. Eklund 2013). For instance, we say that if you make this claim about this case, then you must make the same claim about that case, or else show how they differ. And so on. This important process of comparing cases needn't be impeded by the controversial character of a case.

⁸ As we saw in n.5, some mind-dependent theories try to capture intuitions that favour Mind-Independence via rich views of idealisation, on which ideal agents are committed to certain desires. This won't help this stalemate. If rich views of idealisation are plausible, their defenders are on the same side of the stalemate as Mind-Independence theorists. If they are implausible, then their defenders must retreat to a thin view of idealisation, and are in the same boat as other mind-dependent theories. Either way, this move won't help.

⁹ I'm not saying that *nobody* is persuaded in this way, but I fear that it's uncommon.

So, the problem with controversial cases is relative to our goal of convincing opponents. But you might now worry that I've misunderstood what philosophers aim to achieve when they develop arguments. They're not attempting to convince others to achieve convergence, you may think. They're just trying to find the truth, and we can do that without achieving convergence. So, perhaps we needn't worry about dialectically ineffective arguments, for a philosopher's job is to discover the truth, not to convince others. I think, however, that this is a bad view of philosophical discourse. I agree that philosophers aim to discover what's true and what's not. Or, more carefully, truth is a regulative ideal of philosophical inquiry as it is commonly understood. I also agree that there's more to philosophy than attempts to convince others. But these aims are not mutually exclusive, and my experience is that most philosophical discourse is guided by both. That is, in making arguments we aim to convince others of what we think is the truth. We are rarely so immodest as to expect that our arguments will lead to total convergence, but in trying to convince others of what we think is the truth, convergence lurks in the background as a regulative ideal.

So, even if there are many sound arguments that rely on intuitions about controversial cases, they are still problematic because they fail to persuade opponents. And the problem is big, for The Substantive Strategy that creates it is popular. We thus need to consider how we can overcome the stalemate and move forward. That's a major job, and I can't complete it in this chapter. I can, however, make a start. I diagnose the stalemate's intractability in the next two sections. In later sections I put this diagnosis to work, using it to assess efforts to move things forward via a debunking argument.

6.3 Temperament

Those who spend time arguing about moral reality seem to be thoroughly and persistently divided. In other words, the stalemate discussed above is long-standing and intractable. It is worth considering why this is so, for if we can diagnose how the stalemate becomes entrenched, we'll be in a better position to navigate future metaethical inquiry. My diagnosis of the stalemate's intractability draws on aspects of individual *temperament*. In this section I develop a general view of temperament, to clarify the broad framework within which my diagnosis is located. In the next I develop my diagnosis by focusing on the aspect of temperament that is crucial in this context.

I take temperament to constitute the most significant obstacle to progress in the debate about moral reality. The idea is that, when some aspect of reality is not clear-cut, temperament can play a more decisive role in determining our view of it. So, when we differ in temperament, this leads us to different views of reality. In the case at issue, the effect is that metaethical debates stall. Before explaining how this happens, let's first look at what temperament actually is. The idea that temperament plays a role in philosophical disputes was first put forward by William James, but James developed no explicit definition of what temperament actually is.¹⁰ In fairness, ordinary use doesn't appear to give very sharp boundaries to the term 'temperament.' Providing a fully-fledged *definition* is thus quite hard.

But this doesn't mean that nothing useful can be said. On my view, temperament is broadly speaking a cluster of conative dispositions that mediate an individual's reactions to reality. These conative dispositions may include a desire for simplicity, an expectation of unity, a yearning for explanation, a need for meaning, a bias toward clarity, a hostility to error, a fear of being duped, a wariness of minority views, and so on. Everyone has a temperament, but we also differ in temperament. Differences arise because we can have aspects of temperament more or less strongly. Some have a stronger desire for simplicity than others, some are less wary of minority views, and so on. Thus, whilst we can expect to see fairly broad similarities between individual temperaments, especially within cultures and micro-cultures, we should also expect to find a significant amount of variation.

Note that, being a group of conative *dispositions*, the existence and impact of temperament often remains unconscious. One will typically be unaware of the fact that it mediates one's reactions to reality, for it fixes one's *pre-reflective* sense of how reality will be. It implicitly shapes what we are inclined to see as possible and plausible, and predisposes us to see reality this way or that. True, such predispositions may be overridden by conscious thought. But even then they affect our starting point in inquiry and the assessment of evidence. So, even if we are likely to sometimes be aware of temperament, and can reflectively endorse or veto its influence, its pre-reflective character means that we will usually be unconscious of it.

¹⁰ James (1907). Also relevant are James (1896) and Nietzsche (1886). The account of temperament developed here gels with (and is influenced by) Bordogna (2008: 185).

James' key insight was that temperament plays a role in inquiry. It helps us to identify and assess what counts as evidence for or against a proposition. Relatedly, it influences how one weighs this evidence. It does this by affecting our expectations of, and reactions to, reality. This means that when epistemic agents *differ* in temperament, they will often differ correspondingly in how they count and weigh the available evidence. Of course, differences in temperament will not always lead agents to draw different conclusions about how things really are. Reality often presents itself in fairly clear-cut ways, at least to reasonable and informed people who think about it carefully. When this is the case, there is less space for temperament to fix one's ultimate view. Even if reality conflicts with one's temperament, when things are clear-cut we override this and accept the truth.

At least, this is what conscientious inquirers do. However, reality doesn't always present itself in a clear-cut way, even to conscientious inquirers. Sometimes we face difficult and controversial cases, like those involving ICEs. In such cases, the truth is not clear-cut. When faced with this sort of controversial case, there is therefore room for temperament to more conclusively determine one's ultimate view of the evidence. As a result, temperamental differences may lead us to accept different views of moral reality, by leading us to assess the relevant evidence in different ways.

Now, when arguments about moral reality work in line with The Substantive Strategy, the evidence is taken to be substantive normative intuition. My thought is that temperamental differences may lead us to count and weigh this putative evidence differently, by affecting what we are prepared to say at the substantive normative level. In the next section I explain how this has a lot to do with a specific aspect of temperament. But the lesson of this section is that, whilst a stalemate *arises* because we argue about controversial cases, its *entrenchment* may be due to the fact that the controversy makes room for temperament to more decisively fix our views. A lot more can be said about the nature of temperament itself, but we have yet to assess whether and any how particular differences in temperament do in fact explain the entrenchment of the stalemate. So, rather than exploring the framework of temperament in greater depth, let's see exactly how temperamental differences entrench the stalemate. In the next section I focus on the part of temperament that plays the key explanatory role – the crucial but underdeveloped notion of an *existential need*.

6.4 Existential Needs

By 'existential need' I mean a need for something to be a certain way in order to experience one's life as worthwhile, one's goals as meaningful, or one's choices as non-arbitrary. James did not delineate and discuss the idea of an existential need, but I suggest that they are an important aspect of temperament.¹¹ Our existential needs will often affect our responses to reality, shaping how we count and weigh whatever evidence is available to us.

As with temperamental differences more generally, differences in existential need plausibly have an influence on one's views about some question. For instance, it seems likely that some theists need God to exist to experience their lives as worthwhile. After all, many have been led to despair by the loss of their religious belief. It is plausible that those who have such existential needs will be more inclined to view certain widely accepted facts – the regularity of natural laws, say – as evidence for God's existence. Others lack such existential needs, however, and are thus more likely to discount such facts as evidence for theism, or else will be inclined to give much less evidential weight to them. In this case, differences in existential need result in very different views about a certain aspect of reality.

There are no doubt other debates in which existential needs play a significant role. Some likely candidates include debates about free will, personal identity, and consciousness. But my present claim is local to the Mind-Independence debate. I hold that existential needs help to entrench the stalemate; we can explain persistent differences in intuition about the cases that divide defenders and critics of Mind-Independence by looking at differences in existential need. Some of us may need mind-independent moral truths in order to experience their moral choices as non-arbitrary, for example, whereas others get by with something more metaphysically modest, like idealised desires. This difference in existential need leads us to diverge in what views we are prepared to tolerate or embrace at the substantive moral level. This in turn leads us to different views of moral reality.

¹¹ Something like this notion of existential need plays a role in Nietzschean thought, but I prefer to locate it within a Jamesian framework. It would take me too far afield to explain why I prefer this way of setting things up, but it has to do with the fact that I don't wish to be driven down an ultimately perspectivalist or relativist route by a Nietzschean outlook.

The structure of the explanation is as follows: a stalemate *arises* when we argue about controversial cases, and it gets *entrenched* because the controversy makes room for our existential needs to more decisively determine our ultimate view of moral reality. Some are led to defend to views like Robust Realism, for it is committed to Mind-Independence. But others get pushed down a different path. This is not a matter of vicious wishful thinking, on either side. Instead, virtuous inquirers are led to take a certain view because of how they are disposed to count and weigh the available data. Their existential needs mediate their interpretation of the evidence. So, against what might be taken from an uncharitable reading of James, it's not that we just opt to believe a view because it fits with how we want the universe to be. It's *not* that some believe Mind-Independence because they wish it were true, even though others disbelieve it because they wish differently. Rather, it's that an individual's careful and critical inquiry provides them with apparent evidence for a certain view of moral reality, for this assessment is influenced by temperament and existential need.

Note too that the claim is not that any given argument is rendered *invalid* or *unsound* by virtue of the fact that its author was influenced by existential need. That would be an *ad hominem* charge, and my claims are more subtle. I am suggesting that differences in existential need help to explain why we have persistently differing intuitions about controversial cases involving ICEs. By appealing to differences in existential need, we can enrich our view of why some arguments are not dialectically effective. Because we differ in temperament, we get a persistent clash of intuition and an entrenched stalemate. There's no *ad hominem* here, for the claim concerns the capacity of arguments to convince, not their soundness or validity.

But one might be suspicious of the very idea of an existential need. I can imagine philosophers denying that they are led or misled by such needs. I can also imagine some denying that they have such needs. This is unsurprising. As I've said, the dispositions that constitute temperament often remain unconscious. With existential needs, this is especially likely if these needs are implicitly taken to be satisfied. Just as you tend not strongly experience your need to drink if you're not thirsty, you tend not to strongly experience your existential needs if you assume that reality has what it takes for your moral choices to count as non-arbitrary. It is when existential needs are unsatisfied that their presence is most keenly felt.

Anyway, humans are exactly the type of creature that one would expect to have existential needs.¹² We engage in purposive activity; we pursue goals by making choices. If one judges that reality offers no basis for the value of these goals or the non-arbitrariness of these choices, one might experience it as a rather bleak place. This is likely to have been a source of Tolstoy's period of despair, for example. We don't all have as dramatic a time as Tolstoy, of course. But we all have desires and aims, and we all wonder whether these matter in 'The Grand Scheme of Things.' That is, we can all ask whether reality offers a basis for seeing our goals and decisions as non-arbitrary. This is all it takes to have an existential need.

But how do people differ in their existential needs? One important factor, I suggest, relates to how one feels about the presence of necessity and contingency in ethical life. Some experience a deep and unavoidable link between the contingency of our attitudes and the arbitrariness of any (fundamental) moral norms that constitutively depend on such attitudes. Others, however, don't experience any such link between contingency and arbitrariness. The way an individual feels about this putative link is thus tied up with what they need from reality to experience their moral life as worthwhile and their moral choices as non-arbitrary. Some need a necessary and thus mind-independent moral reality for their existential needs to be satisfied, whereas contingent desires and aims are enough for others.

For instance, a recent interview with Derek Parfit describes him as thinking that, without objective moral truth, the world "would be a bleak place in which nothing mattered." This thought "horrifies" Parfit, who goes on to say that in such a world we would have "no reasons to try to decide how to live. Such decisions would be arbitrary ... We would act only on our instincts and desires, living as other animals live." However, Parfit also recognises that many of his colleagues "not only do not believe in objective moral truth – they don't even find its absence disturbing." In short, Parfit seems troubled by the idea that moral truth is as contingent as many metaethicists suppose. His existential needs are unsatisfied by views that allow moral normativity to depend on contingent desires and aims.¹³

¹² For relevant discussion, see Cooper (2002), Ratcliffe (2005) and Kidd (2013).

¹³ The quotes from or about Parfit in this paragraph are from an interview conducted by MacFarquhar (2011). Similar thoughts are expressed and defended in Parfit (2011).

Other philosophers see things differently. James Lenman answers Parfit's concerns directly:

The ... realism favoured by Parfit is one of a number of philosophical positions where it is very common for those who believe them not only to consider them intellectually compelling but to think it would be catastrophically awful were they wrong. But I don't get it (2009: 37).

Lenman doesn't get why the contingency of our desires should make any reasons that derive from them arbitrary. One's desires don't matter from a 'zoomed out' standpoint, but this is not our standpoint.¹⁴ Street is also untroubled by the 'ultimate' contingency of reasons, comparing it to the contingency of love.¹⁵ She notes that there are many people with whom one could have fallen in love, and we can recognise this without it thereby undermining "one's lifelong love and commitment to the person whom one did, as it so happened, actually meet ... and build a happy life with" (2012: 57). Street sees moral contingency as equally untroubling. Indeed, one might even embrace these contingencies as liberating. They allow us to construct our own paths without the constraint of some independent authority.

In short, there are people like Lenman and Street who have more relaxed existential needs than people like Parfit, and this relates to their feelings about contingency and arbitrariness. I don't mean to single out these philosophers. People in general are temperamentally led to count and weigh the substantive normative evidence differently, and existential needs do explanatory work when we ask why this is the case. Moreover, it is not clear who must 'veto' their existential needs, for we're dealing with controversial cases. Both sides are led to stalemate by temperament, but it is not clear who is being *mised*. We thus need further argument to see who should override their existential needs. And we must go beyond The Substantive Strategy to achieve this. More broadly, to move forward we must go beyond arguments in which temperament inevitably plays a role. Otherwise stalemate will re-emerge; we'll only be moving the bump in the rug. In §6.5 to §6.7 I consider an increasingly influential style of argument, one that may seem to move the debate forward against Robust Realism.

¹⁴ Lenman (2009: 36-38). Cf. Hare (1972) and Tiberius (2012).

¹⁵ Cf. Lenman (1999).

6.5 The Debunking Strategy

How, then, can we move forward and make progress in the debate about moral reality? By identifying styles of argument that enable us to avoid or overcome the problematic impact of temperament and existential need. Unfortunately, there is not space here to identify and assess every such style of argument. I will therefore focus on one that may seem especially promising, that has been highly influential in recent years, and that often comes up in critiques of Robust Realism. This is ‘The Debunking Strategy,’ and it is designed to threaten the commitment to Mind-Independence.¹⁶ In this section, I outline the heart of this way of trying to move the debate forward. In the next two sections, I argue that it fails to either avoid or overcome temperament. Divisive temperamental factors play a role even within debunking arguments, and the stalemate re-emerges in a different way.

The Debunking Strategy has been developed in various ways, but the basic idea is that we can undercut the intuitions that seem to support Mind-Independence by looking at their causal history. Recent versions of The Debunking Strategy have typically focused on the evolutionary origin of our moral attitudes, and this sort of evolutionary debunking argument will be my focus here.¹⁷ The idea is that, by explaining our moral views in evolutionary terms, we expose that it is a mistake to invoke something mind-independent to vindicate those views. If this evolutionary claim is correct, it will be a serious problem for Mind-Independence and Robust Realism.

¹⁶ A way of trying to move the debate forward in favour of Mind-Independence is to appeal to ‘companions in innocence’ from other normative domains (like the epistemic domain) in order to vindicate mind-independent entities in the moral domain. For this strategy, see Hampton (1998), Stratton-Lake (2002: xxv-xxvi), Shafer-Landau (2003: 113), Cuneo (2007a), Bedke (2010), and Rowland (2013). For critiques, see Heathwood (2009), Cowie (2014b), and Ingram (MS). Another way of trying to move the debate forward in favour of Mind-Independence is to argue that the idealisation required to make mind-dependent theories come close to extensional adequacy cannot be motivated in a principled way. For this strategy, see Enoch (2005). For a response, see Sobel (2009). Other argumentative strategies are also available. For instance, we might look to experimental philosophy or linguistics to improve our view of the substantive moral data. I hope to explore various efforts to move the debate forward in future, to see if they succeed in avoiding or overcoming temperament. My suspicion is that they won’t, though I am a pessimist.

¹⁷ For a non-evolutionary version of The Debunking Strategy, see Nietzsche (1887). For evolutionary versions, see Ruse (1986), Joyce (2006), Street (2006), and Kitcher (2011).

There are two core ideas that motivate The Debunking Strategy. The first is that we have our moral views because we in some way evolved to have them. The second is that appreciation of this evolutionary history makes trouble for Mind-Independence, by exposing how our moral views are sensitive to the evolutionary ‘aims’ of survival and reproduction, and not to a mind-independent moral reality.¹⁸ How can we develop these two basic thoughts into an interesting argument? Let’s start with the idea that evolution influenced our moral views. Debunkers allow that many things have had such an influence – culture, rational reflection, and so on – but they say that evolution was especially significant. To put it a little crudely, a disposition to judge that it is immoral to kill babies would be much more likely to promote the survival and reproduction of one’s genes than a disposition to judge that it is morally permissible to kill babies. The forces of evolution – natural selection, sexual selection, and so on – will thus have favoured those with the former disposition. So it’s no surprise that we judge killing babies immoral. Of course, evolution didn’t directly select all of our specific moral views. The thought is that it does so indirectly, by selecting broad evaluative tendencies that then affect our specific moral views.

The first premise in a debunking argument is thus something like the following:

- (1) Our substantive moral views are influenced significantly by our evolutionary history.

Note that (1) is an empirical claim. It looks plausible, but a strong reading of it isn’t easy to secure. Even detailed discussions come with caveats. For instance, Richard Joyce’s answer to the question of how our moral views evolved is “provisional and to a degree speculative, since the present evidence does not warrant answering the question in either a positive or a negative way with any confidence” (2006: 2). Similarly, Philip Kitcher says that reconstructing “the *actual* history of the ethical project, from its beginnings to the present, is plainly beyond the evidence available – and probably beyond the evidence anyone could ever hope to obtain” (2011: 11). This will become more important later, as we consider how different ways of reading (1) can affect the viability of debunking arguments. But what matters now is that something like (1) is any such argument’s first premise.

¹⁸ Talk of evolutionary ‘aims’ is metaphorical, but useful for my purposes.

The second step in the debunking argument can be summarised as follows:

- (2) So, we have no good reason to believe our substantive moral views to be mind-independently true.

This is what many debunkers take to be the key lesson that we can draw from (1). To move us from (1) to (2), they observe that we must somehow explain the link between our moral views on the one hand and evolution on the other. If we take our moral views to reflect a mind-independent moral reality, we might explain this by saying that we evolved the ability to reliably ‘track’ this part of reality. Just as cheetahs were selected for speed, humans were selected to discern mind-independent normative truths. This is Parfit’s view.¹⁹ Street (2006) calls it ‘The Tracking Account,’ and argues that the explanation offered by this account is not as good as her own. According to what Street calls ‘The Adaptive Link Account,’ we have the moral views we do because they helped to form adaptive links between our ancestors’ circumstances and advantageous responses to those circumstances. For instance, those who tended to experience caring for offspring as *required* tended to have more reproductive success as a result.

Street argues that The Adaptive Link Account is simpler, clearer, and more illuminating than The Tracking Account. And this allows the move from (1) to (2) – we have no good reason to believe that our moral views reflect something mind-independent, because the best explanation of evolution’s influence on such views is one that makes no reference to such a reality.²⁰ Debunkers might hope to undercut Mind-Independence from (1) and (2) alone, but that would be too quick. Katia Vavova (2014) points out that we’ll first need something like the following principle:

- (3) If you have no good reason to believe that p , then you cannot rationally maintain that p .²¹

As we’ll see later on, this principle leads to problems for debunkers. But what matters for now is that it helps get them to the conclusion that they want.

¹⁹ For details, see Parfit (2011 v.2: §119).

²⁰ Of course, it may still be that a mind-dependent view of moral truth can be given.

²¹ This is different to Vavova’s phrasing, but not in a way that matters. This principle is endorsed by Street (2015). For related points, see Vavova (2015).

Let's bring these points together to clarify how this influential version of The Debunking Strategy is used against Mind-Independence:

- (1) Our substantive moral views are influenced significantly by our evolutionary history.
- (2) So, we have no good reason to believe our substantive moral views to be mind-independently true.
- (3) If one has no good reason to believe that p , then one cannot rationally maintain that p .
- (4) So, we cannot rationally maintain that our substantive moral views are mind-independently true.

'One cannot rationally maintain that p ' is more modest than 'not- p ', so (4) doesn't deny Mind-Independence. This modesty is needed as, clearly, one cannot infer not- p from the fact that one has no good reason to believe that p . Still, the conclusion makes a strong claim. It is a serious problem for Mind-Independence, for it doesn't rely on divisive moral views. In fact, temperament doesn't *seem* to play any role here, for the argument centres on empirical claims about how we evolved, not on any substantive moral intuition.

6.6 Good Reason

This debunking argument may not rely on substantive moral intuitions, but it does rely on a substantive *epistemic* intuition. This is because it relies on (3), which makes a substantive epistemic claim about when one can rationally maintain a belief. So we must ask whether (3) is intuitively correct.

One might well doubt this, for (3) is not the only game in town. To see this, consider the contrast the following two principles:

- (3) If one has no good reason to believe that p , then one cannot rationally maintain that p .
- (3*) If one has good reason to believe that not- p , then one cannot rationally maintain that p .

Vavova (2014: 85) notes that (3) and (3*) look similar, but are crucially distinct. (3) treats beliefs as guilty until proven innocent. It doesn't allow you to rationally maintain that p until you have a good reason for thinking it true. (3*) treats beliefs as innocent until proven guilty. It allows that p might be rationally maintained until one has a good reason to think that not- p .

Which of these epistemic norms is intuitively more compelling? Vavova argues that (3) is intuitively problematic, for it risks proving too much. To see this, consider first that debunkers must see 'good reason' as referring to a reason that is independent of our substantive normative views. They have to say this to avoid responses like those offered by Erik Wielenberg (2010), David Enoch (2010b), and Knut O. Skarsaune (2011) who try to explain the non-accidental correlation of adaptive moral views and the mind-independent moral facts by relying on substantive moral claims.

For instance, Enoch uses the claim that 'survival is good' to show how correlations may arise. If survival is mind-independently good, then moral views that evolved to 'aim' at survival will at least roughly match the mind-independent moral facts. By showing how the 'aim' of survival led to moral belief *M*, we thus support rather than debunk *M*: survival is mind-independently good, *M* helps survival, so – all else being equal – *M* is mind-independently good. Debunkers respond by saying that this begs the question; our reason to maintain that morality's mind-independence is not debunked cannot itself be or rest on any mind-independent moral view, for it is precisely such views that are in question. So, in short, a 'good' reason to think this has to be independent of everything that is in question.

But if this is the case, then (3) proves too much.²² Vavova argues that, if we make this independence claim, we have no good reason to think that sensory perception leads us to true beliefs:

[If] we set aside all that is in question, we must set aside all beliefs gained by perception. This includes all scientific beliefs, like the belief that evolutionary theory is true. Without those, we cannot evaluate the rationality of beliefs formed by perception. We can test the reliability of a particular sense modality by granting the reliability of others ... But if we cannot rely on our senses, we have nothing with which to evaluate reliability. We have set aside too much (2014: 82-83).

The result is that our justifications will run out, and our beliefs will rest on nothing. Operating with (3) thus means that The Debunking Strategy collapses into general scepticism. This is a result that debunkers wish to avoid.

²² Cf. Shafer-Landau (2012: 12-13).

There are two ways for them to go at this point. The first is to stick with something much like (3), whilst introducing a principled revision so that it stays local to the moral domain and avoids a more general sceptical challenge. The second is to replace (3) with (3*), and revise the rest of the argument accordingly. I don't think that either move works, for in the end both fail to avoid or overcome temperament. But let's consider them in turn.

6.7 The Temperament in Debunking

How might debunkers seek to revise (3) so that it stays local to the moral domain? The most principled revision is likely to be the following:

(3-sui-generis) If one has no good reason to believe that p ,
where p commits us to a *sui generis* property,
then one cannot rationally maintain that p .

The idea is that we must be wary of accepting *sui generis* properties. After all, such properties are unique and thus discontinuous with other aspects of reality. They create a less simple and unified picture, and can thus seem awkward additions to an ontology. So maybe we should treat beliefs that commit us to them as guilty until proven innocent. Of course, most beliefs don't commit us to *sui generis* properties. But if moral beliefs do, as many have thought, then (3-sui-generis) allows a revised debunking argument to go through without its overgeneralising. This revised argument could make problems for our moral beliefs whilst leaving all of our other beliefs intact.²³

My worry about this idea is that there is room for temperamental division about (3-sui-generis). To see this, note that it is motivated by the appeal of simplicity and unity. The idea is that we should be suspicious of properties that are discontinuous with the rest of reality, for they create a more complex and disunified view. But we differ in the extent to which we share this suspicion. We differ in how attractive we find simplicity and unity, and in how we weigh these against a need to save the appearances. And this is a matter of temperament. Some are temperamentally inclined to save all of the appearances, even if it entails a more complex and less unified view of reality. Others, however, are willing to save most and not all of the appearances, if it will mean a simpler and more unified view of reality.

²³ Note that a revised argument will need an additional premise stating that moral beliefs do in fact commit us to *sui generis* moral properties.

(3-sui-generis) might well seem an attractive epistemic norm to those who are temperamentally inclined to accept a simpler theory that fails to capture some of the appearances, but it will be far less attractive to those who are temperamentally inclined to accept the more complex theory that can deliver all the appearances. So, if a revised version of the debunking argument were to rely on this principle, it still wouldn't move the debate forward. A stalemate would re-emerge at a different level, for some will be inclined to reject (3-sui-generis) and others will be inclined to accept it, and this divergence in substantive epistemic intuition is liable to become entrenched because it is linked to temperamental differences. So, this is not the best way for a debunker to go. I can't think of a similarly principled revision to (3) that will do better, so let's consider the second move.

The second move is to reject (3) and its variants altogether, and to move instead to (3*). If we make this move, we will also need to change the second premise in the debunking argument. The revised argument would go as follows:

- (1) Our substantive moral views are influenced significantly by our evolutionary history.
- (2*) So, we have good reason to believe that our substantive moral views are not mind-independently true.
- (3*) If one has good reason to believe that not-*p*, then one cannot rationally maintain that *p*.
- (4) So, we cannot rationally maintain that our substantive moral views are mind-independently true.

This argument probably provides the most promising way to develop The Debunking Strategy. But it is also problematic, for now the debunker has to do something rather tricky. Specifically, they have to show that we have good reason to doubt that our moral views are mind-independently true.

This is not easy. A move from (1) to (2*) is hard to pull off, for (2*) is much stronger than (2). To deliver (2*) we need a strong reading of (1). We need to demonstrate that, directly or indirectly, evolutionary forces had by far the most significant impact on our moral views, for modest readings of (1) – according to which evolutionary forces are an influence, but not the only and not the biggest – won't permit the move to (2*). The trouble is that the strong view of (1) is not obviously superior to a modest view.

Consider a modest reading of (1). William FitzPatrick argues that, whilst evolution gave us capacities to reason in various fields, applying these to specific and sophisticated matters is something that happens in “cultural contexts, through relevant forms of training within traditions of inquiry into the subject matter in question” (2015: 887). Our cognitive capacities might thus non-accidentally track truths in some domain, even if the materials for their development evolved.²⁴ And this is in fact how various areas of inquiry work: knowledge of complex mathematical and metaphysical facts requires training in a cultural context, even if the raw materials required for reasoning about mathematical and metaphysical questions came from evolutionary forces. A similar claim can be made for ethics. On this modest approach, you can agree with (1) without thinking that it implies (2*). For on this view, evolution has an impact insofar as it provided us with the materials needed for moral or normative reasoning, but moral and normative views develop within culture, and it is here that our capacity to reason is trained to discern mind-independent facts in this area of inquiry. So, a modest view of (1) is evidently not going to entail (2*).

We can now ask whether the modest view of (1) is obviously less convincing than the ambitious view. Recall that the ambitious view tells us that, directly or indirectly, evolution had by far the biggest impact on our normative views. This is a strong empirical claim, and it will be hard to get the evidence that would be required to vindicate it. As we’ve seen, even Joyce and Kitcher observe that the relevant evolutionary claims are speculative. The strong view of (1) has yet to be secured, so it is a gamble to rely on it to deliver (2*). To be clear, I’m not saying that the strong view is false, and I’m not saying that the modest view suggested by FitzPatrick is true. I’m just saying that the fact of the matter is not now clear. And it may never become clear, given that (as Kitcher says) we’ll probably never discover the empirical evidence that will be required to confidently and accurately reconstruct the actual history of our substantive normative views.

²⁴ We must take care. There’s evidence that norm-thought is not a cultural application of highly general rational capacities, but is itself evolved (Sripada and Stich 2006). I accept this, but it doesn’t conflict with my point. Even if evolution gave us capacities for specific domains, these are developed within cultures, and are applied in more subtle ways than they will have been by our ancestors. They are also applied in concert with other rational capacities, often in (semi-)public debates. This gives an empirically kosher view of (1) that is more modest even than The Adaptive Link Account. Thanks to Alex Duval for discussion.

So, we have at least two live options when it comes to interpreting the first premise in the debunking argument. One of these permits a move to (2*), but the other doesn't. It may be that time will tell in the debunker's favour, or it may be that time will tell in favour of Mind-Independence. In the meantime, we can plausibly suppose that what divides defenders and critics of The Debunking Strategy are temperament and existential need. Given that the fact of the matter is not at present clear-cut, there is space for temperamental factors to play a bigger role in shaping our ultimate views of this question. Whilst some are temperamentally inclined to think that the evidence favours the more ambitious interpretation of (1), others are temperamentally inclined to think that it favours the more modest interpretation. We are again divided by our temperaments and existential needs.

Of course, even though we may never get the empirical evidence required to decide between the two readings of (1), it is also possible that we will. It would thus be premature to claim that The Debunking Strategy cannot succeed in moving things forward. Still, what we've seen in this section is that it has yet to move things forward, and that it might never do so. So we can at least say that a belief in Mind-Independence has yet to be undercut by The Debunking Strategy. The stalemate just re-emerges at another level, and we should be open to the possibility that it won't go away.

6.8 The Limits of Metaethics

There are other ways of trying to move the debate forward. In §2.8 I said that Robust Realism should be defended primarily by (non-substantive) metaphysical argument. Maybe a scrupulously metaphysical approach, centring on the development and comparison of ways moral reality *might* be, can avoid divisive temperamental factors. I hope so, but even here I'm not sure. A metaphysical approach seems more promising than the moral approach, but I suspect that it will still be hard to overcome temperament and move forward. I don't suggest that it will be impossible. Now that we have an idea of what holds us back, we are in a far better position to see how we can move forward. Yet we have just seen reason to be pessimistic about one of the most influential attempts to do this, so it is worth asking what we should do if all efforts fail. Not because we are prematurely concluding that progress is unachievable, but because it may bring some comfort to know what our situation will be like if progress is not in fact achieved.

Suppose, then, that there is no way out of the stalemate. Can we continue making arguments, even though we know that they won't be dialectically effective for those who differ in temperament? What would metaethical discourse be like if we stopped the attempt to convince one another of what we take to be the truth about moral reality? Would there still be life in our debates? Can we carry on with more limited ambitions? I think so. For one thing, even if we cannot convince opponents of what we take to be the truth, we can do other stuff. I noted in §6.2, for example, that difficult cases can be used to preach to the choir, and can expose and elucidate a theory's commitments. These activities are interesting and philosophically important, even if they rarely involve changing anyone's view.

The more challenging question is whether it can be intellectually acceptable to believe a theory when you know that at least some of your epistemic peers persistently disagree with you about it. Some think that, in such circumstances, the rational thing to do is to suspend judgement.²⁵ Perhaps this is sometimes the case – for instance, when there is a good chance that further investigation will clear things up – but it is less attractive in this situation. We're assuming that various attempts to move metaethical debate forward have failed, and that we have seen how they fail to avoid or overcome the influence of temperament. If we suspend judgement, we'll probably never be entitled to have any metaethical view at all. And that seems unattractive, given that – as we've seen – theories of moral reality can often be important to the satisfaction of our existential needs.

A good way to go is to allow that there are circumstances in which it is intellectually acceptable to believe some view, even when one's belief is ultimately due more to temperament than argument. There isn't space to develop this possibility in detail here, for it will inevitably raise various interesting and difficult epistemological questions. Nevertheless, an idea that I find attractive involves emphasising the importance of humility in intellectual life. If we are appropriately guided by the virtue of intellectual humility, there is plausibly scope for an agent to believe the account of normative reasons to which they are temperamentally inclined. This is so even if they are unable to give a dialectically effective argument for their view.

²⁵ Sidgwick (1907: 341), Kelly (2006), Elga (2007), McGrath (2008), and Crisp (2011).

The notion of humility that matters here is one on which it is opposed to hubris on the one hand and a lack of self-confidence on the other.²⁶ Roughly, humility is the virtue that regulates confidence. As an epistemic virtue, it helps epistemic agents to avoid serious excesses and deficiencies of confidence in epistemic activities like forming beliefs and making arguments.²⁷ Individuals with the virtue of humility understand the strengths and limitations of their own cognitive capacities, and those of humanity in general. They form beliefs with an appropriate degree of confidence, and argue with a recognition of what argument can and can't achieve.

It is worth emphasising that, on this view, being humble – that is, having the virtue of humility – is not a matter of *lacking* confidence. In fact, humble people can be very confident about the truth of a proposition and the effectiveness of an argument for it – when it is clear that *p*, they will be very confident that *p*. Not excessively confident, of course, but not deficient in confidence either. They'll strike a suitable balance. However, when the matter is less clear-cut, humble people might form a belief with a lower level of confidence, or they might choose to suspend judgement if that will help. It depends on the situation, but we're discussing a situation in which the truth is not *just* unclear. In addition to being unclear, it is also the case that our efforts to find the truth are unavoidably influenced by temperament and existential need. What is the humble agent to do in this situation?

When temperament's impact is unavoidable, so that suspending judgement is pointless, this shouldn't stop humble people from forming beliefs. Especially when the issue is so important that abandoning inquiry into it would adversely affect their lives. When an issue has great practical significance, humble inquirers needn't ignore it or stay neutral. They can engage in inquiry about it, and form beliefs about its subject-matter. But their beliefs will be held with appropriate levels of confidence. This might be quite low, given the effect that temperament is likely to be having on their formation of that belief. Indeed, it might be such a low degree of confidence that it doesn't extend far beyond the belief that this is a way things *might* be. But the belief is one that they are permitted, humbly, to form.

²⁶ Humility may oppose various vices, including hubris, vanity, and pretentiousness. For more on this, see Roberts and Wood (2006: Ch. 9).

²⁷ For detailed discussion, see Kidd (forthcoming).

In the specific case at issue, we are therefore permitted to argue about moral reality, and to form beliefs about the relation between moral properties and our attitudes. This is so even if such beliefs are often due more to temperament than argument, and even if there is little or no hope of overcoming temperament. But we must cultivate humility, so that we grasp the level of confidence that is appropriate for these beliefs about moral reality. This means that, when people come to different views of moral reality, they can rationally maintain their views as long as they are appropriately humble about them.²⁸ I suppose that some people might be disappointed by the modesty of this approach to metaethical inquiry. Others will no doubt see it as reassuringly self-aware. Personally, I tend to identify more with the latter camp. But perhaps this too is a matter of temperament.

6.9 Conclusion

It is too soon to say whether we will eventually overcome the stalemate that exists in much of contemporary metaethics. But to do so we will have to come up with argumentative strategies that can avoid or overcome the influence of temperament. Although the stalemate arises because we tend to argue about hard cases – cases about which reasonable and thoughtful can persistently disagree – it becomes entrenched because temperament and existential need step in and influence our view of moral reality. I have given reasons to be pessimistic about a prominent style of argument that might seem to move the debate forward at the expense of Robust Realism and Mind-Independence. I don't yet know whether, in the future, we will find other ways of arguing that help us to make progress. But if we can't, that's not the end of the world. It simply means that we need to be more humble.

²⁸ One might worry that humble beliefs won't fully satisfy existential needs. This might be right – they won't do as much as beliefs that we hold confidently. But we're not entitled to high levels of confidence in the situation at issue. And, even if diminishing confidence will augment existential dissatisfaction, it won't do this nearly as much as the abandonment of inquiry or the suspension of judgement. Moreover, *some* existential dissatisfaction may be just what is called for in the dialectical situation under consideration. Thanks to Neil Williams for pressing me on this point.

7 Conclusion

I am a Robust Realist, and I hope that I am humble about it. I try to be, at any rate. Being morally and intellectually virtuous is a lot easier said than done.¹

Humility is especially needed in my case, for there are important issues that I have set aside. My primary aim has been to provide Robust Realism with a compelling metaphysics. I thus paid little attention to the epistemic, semantic, and psychological commitments that are associated with this view. There is a reason for this. I said at the start that Robust Realism is an *essentially* metaphysical view. The non-metaphysical claims that are associated with it are important and interesting, but they are not definitive of Robust Realism. Even so, providing a robustly realistic view of moral knowledge, language, and psychology is not something that can be set aside forever. I'll leave this as a task to be undertaken in the future. Until it's completed I must make a special effort to be humble about moral reality.

So there's some stuff that I haven't yet done, and it is stuff that I'll have to do in the future. But there's also a fair bit that I have done. I have argued that it is possible to read claims about moral reality in a robustly metaphysical sense. Whilst a substantive moral reading is available, it is not the only available reading. I have also argued that the commitment to irreducible moral properties is not undermined by the supervenience of moral properties on descriptive properties, at least not if we claim that there are necessary moral truths. I developed a model of necessary moral truths as norms that have had contingent limiting factors stripped away, and I have argued that the categoricity of moral norms makes it plausible to see moral reality as existing beyond the purview of natural and social science.

The metaphysical views that I have developed for Robust Realists have some motivation, and should be taken seriously. But to say that they should be taken seriously is not to say that they are true. I haven't tried to show that Robust Realism is actually true, and I don't know if we'll ever find a dialectically effective argument for it given the pervasive impact of temperament. Maybe we'll discover such an argument one day. Maybe we won't.

¹ Back in §2.6, I pointed out that philosophy PhD programmes are not production lines for moral saints. Well, they're not production lines for epistemic saints either.

This suggests even more work that must be done in the future. We need to identify further ways of arguing about moral reality, and we need to assess whether they succeed in avoiding or overcoming temperament. I tend to be pessimistic about these things. I'm not certain, but I strongly suspect that avoiding or overcoming temperament will be so difficult that it is unlikely to be achieved by human epistemic agents. And even if we do succeed in avoiding or overcoming temperament, it's not as if this will ensure that the nature of moral reality will immediately become clear to us.

That doesn't mean we should stop thinking about moral reality, of course. I'd be delighted to see arguments that can move things forward, and that therefore prove my current pessimism to have been ill-founded. Of course, even if we do succeed in moving the debate forward, we'll still have to be humble in our metaethical beliefs. After all, humility is just the virtue that regulates confidence. It ensures that we are not excessively or deficiently confident in epistemic affairs, and it's possible that the humble metaethicists of the future will be permitted to believe things about moral reality with a great deal more confidence than the humble metaethicists of the present. To settle this issue, we'll just have to carry on talking about moral reality. We'll have to seek new ways of arguing, in order to evaluate their ability to avoid or overcome temperament. We'll have to make our current theories more refined, in order to put our views on a more secure footing. And we'll have to identify new types of theory, in order to assess whether any of them have greater success than those that are currently available to us. So, there is a great deal more to say. But you have to stop somewhere.

References

- Anscombe, G.E.M. 1958. Modern Moral Philosophy. *Philosophy* 33, 1-19.
- Armstrong, D.M. 1989. *A Combinatorial Theory of Possibility*. New York: Cambridge University Press.
- 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Audi, R. 2004. *The Good in the Right: A Theory of Intuition and Intrinsic Value*. Princeton: Princeton University Press.
- 2013. *Moral Perception*. Princeton: Princeton University Press.
- Ayer, A.J. 1936. *Language, Truth and Logic*. London: Penguin, 2001.
- Bedke, M. 2010. Might All Normativity Be Queer? *Australasian Journal of Philosophy* 88, 41-58.
- Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.
- 1993a. How to Be an Ethical Anti-Realist. In *Essays on Quasi-Realism*. Oxford: Oxford University Press, 166-181.
- 1993b. Moral Realism. In *Essays in Quasi-Realism*. Oxford: Oxford University Press, 111-129.
- 1993c. Supervenience Revisited. In *Essays in Quasi-Realism*. Oxford: Oxford University Press, 130-148.
- 1996. Securing the Nots: Moral Epistemology for the Quasi-Realist. In W. Sinnott-Armstrong and M. Timmons (eds.), *Moral Knowledge? New Readings in Moral Epistemology*. New York: Oxford University Press, 82-100.
- 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- 2010a. Truth, Beauty and Goodness. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 5*. Oxford: Oxford University Press, 295-314.
- 2010b. The Majesty of Reason. *Philosophy* 85, 5-27.
- Bloomfield, P. 2009. Archimedeanism and Why Metaethics Matters. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 4*. Oxford: Oxford University Press, 283-302.

- Bordogna, F. 2008. *William James at the Boundaries: Philosophy, Science, and the Geography of Knowledge*. Chicago: University of Chicago Press.
- Boyd, R. 1988. How to be a Moral Realist. In G. Sayre-McCord (ed.), *Essays on Moral Realism*. New York: Cornell University Press, 181-228.
- 2003a. Finite Beings, Finite Goods: The Semantics, Metaphysics, and Ethics of Naturalist Consequentialism, Part I. *Philosophy and Phenomenological Research* 66, 505-553.
- 2003b. Finite Beings, Finite Goods: The Semantics, Metaphysics, and Ethics of Naturalist Consequentialism, Part II. *Philosophy and Phenomenological Research* 67, 24-47.
- Brink, D. 1984. Moral Realism and the Sceptical Arguments from Disagreement and Queerness. *Australasian Journal of Philosophy* 62, 111-125.
- 1989. *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press.
- 2001. Realism, Naturalism, and Moral Semantics. In E.F. Paul, F.D. Miller, Jr. and J. Paul (eds.), *Moral Knowledge*. Cambridge: Cambridge University Press, 154-176.
- Brown, C. 2011. A New and Improved Supervenience Argument for Ethical Descriptivism. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 6*. Oxford: Oxford University Press, 205-218.
- 2012. Still No Redundant Properties: Reply to Wielenberg. *Journal of Ethics and Social Philosophy* (<http://www.jesp.org>), November 2012.
- Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.
- Cohen, G.A. 1996. Reason, Humanity, and the Moral Law. In C.M. Korsgaard, *The Sources of Normativity*. Cambridge: Cambridge University Press, 167-188.
- Colyvan, M. 2010. There is No Easy Road to Nominalism. *Mind* 119, 285-306.
- Cooper, D.E. *The Measure of Things: Humanism, Humility and Mystery*. Oxford: Oxford University Press.

- Copp, D. 2003. Why Naturalism? *Ethical Theory and Moral Practice* 7, 179-200.
- Cowie, C. 2014a. A New Explanatory Challenge for Non-Naturalists. *Res Philosophica* 91, 661-679.
- 2014b. Why Companions in Guilt Arguments Won't Work. *Philosophical Quarterly* 64, 407-422.
- Crisp, R. 2011. Reasonable Disagreement: Sidgwick's Principle and Audi's Intuitionism. In J. Graper Hernandez (ed.), *The New Intuitionism*. London: Continuum, 151-168.
- Cullison, A. 2009. Three Millian Ways to Resolve Open Questions. *Journal of Ethics and Social Philosophy* 3, 1-17.
- Cuneo, T. 2006. Saying What We Mean: An Argument Against Expressivism. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 1*. Oxford: Oxford University Press, 35-71.
- 2007a. *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- 2007b. Recent Faces of Moral Nonnaturalism. *Philosophy Compass* 2, 850-879.
- 2014. *Speech and Morality: On the Metaethical Implications of Speaking*. Oxford: Oxford University Press.
- Cuneo, T. and Shafer-Landau, R. 2014. The Moral Fixed Points: New Directions for Moral Non-Naturalism. *Philosophical Studies* 171, 399-443.
- Dancy, J. 1993. *Moral Reasons*. Oxford: Blackwell.
- 2004a. *Ethics Without Principles*. Oxford: Oxford University Press.
- 2004b. On the Importance of Making Things Right. *Ratio* 17, 229-237.
- Darwall, S, Gibbard, A, and Railton, P. 1992. Toward fin de siècle ethics: Some Trends. *Philosophical Review* 101, 115-189.
- Dreier, J. 1992. The Supervenience Argument against Moral Realism. *The Southern Journal of Philosophy* 30, 13-38.
- 2002. Metaethics and Normative Commitment. *Philosophical Issues* 12, 241-263.

- 2004. Meta-ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives* 18, 23-44.
- 2015. Explaining the Quasi-Real. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 10*. Oxford: Oxford University Press, 273-297.
- Dworkin, R. 1996. Objectivity and Truth: You'd Better Believe It. *Philosophy and Public Affairs* 25, 87-139.
- 2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
- Ehrenberg, K. 2008. Archimedean Metaethics Defended. *Metaphilosophy* 39, 508-529.
- Eklund, M. 2013. Trends and Progress in Philosophy. *Metaphilosophy* 44, 276-292.
- Elga, A. 2007. Reflection and Disagreement. *Noûs* 61, 478-502.
- Elster, J. 2011. How Outlandish Can Imaginary Cases Be? *Journal of Applied Philosophy* 28, 241-258.
- Enoch, D. 2005. Why Idealize? *Ethics* 115, 759-787.
- 2007. An Outline of an argument for Robust Metanormative Realism. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 2*. Oxford: Oxford University Press, 21-50.
- 2010a. How Objectivity Matters. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 5*. Oxford: Oxford University Press, 111-152.
- 2010b. The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With It. *Philosophical Studies*, 148, 413-438.
- 2011. *Taking Morality Seriously: A Defence of Robust Realism*. Oxford: Oxford University Press.
- Field, H. 1980. *Science Without Numbers: A Defence of Nominalism*. Oxford: Blackwell.
- Fine, K. 2002. The Varieties of Necessity. In T.S. Gendler and J. Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Oxford University Press, 253-281.

- Finlay, S. 2008. The Error in the Error Theory. *Australasian Journal of Philosophy* 89, 519-534.
- 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- FitzPatrick, W.J. 2008. Robust Ethical Realism, Non-Naturalism, and Normativity. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 3*. Oxford: Oxford University Press, 159-205.
- 2011. Ethical Non-Naturalism and Normative Properties. In M. Brady (ed.), *New Waves in Metaethics*. New York: Palgrave MacMillan, 7-35.
- 2014. Skepticism about Naturalising Normativity: In Defence of Ethical Nonnaturalism. *Res Philosophica* 91, 559-558.
- 2015. Debunking Evolutionary Debunking of Ethical Realism. *Philosophical Studies* 172, 883-904.
- Foot, P. 1972. Morality as a System of Hypothetical Imperatives. *Philosophical Review* 81, 305-316.
- 2001. *Natural Goodness*. Oxford: Oxford University Press.
- Fricker, M. 2016. What's the Point of Blame: A Paradigm Based Explanation. *Noûs* 50, 165-183.
- Garner, R. 1990. On the Genuine Queerness of Moral Properties and Facts. *Australasian Journal of Philosophy* 68, 137-146.
- 2007. Abolishing Morality. *Ethical Theory and Moral Practice* 10, 137-146.
- Gauthier, D. 1986. *Morals by Agreement*. Oxford: Oxford University Press.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.
- 1999. Morality as Consistency in Living: Korsgaard's Kantian Lectures. *Ethics* 110, 140-164.
- 2002. Knowing What to Do, Seeing What to Do. In P. Stratton-Lake (ed.), *Ethical Intuitionism: Re-evaluations*. Oxford: Oxford University Press, 212-228.
- 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

- 2011. How Much Realism? Evolved Thinkers and Normative Concepts. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 3*. Oxford: Oxford University Press, 33-51.
- Goldman, A.H. 2009. *Reasons from Within: Desires and Values*. Oxford: Oxford University Press.
- Hale, B. 2013. *Necessary Being: An Essay on Ontology, Modality, and the Relations Between Them*. Oxford: Oxford University Press.
- Hampton, J. 1998. *The Authority of Reason*. Cambridge: Cambridge University Press.
- Hare, R.M. 1952. *The Language of Morals*. Oxford: Oxford University Press.
- 1972. Nothing Matters. In *Applications of Moral Philosophy*. London: MacMillan.
- 1979. What is Wrong with Slavery? *Philosophy and Public Affairs* 8, 103-121.
- 1981. *Moral Thinking: Its Levels, Method, and Point*. Oxford: Oxford University Press.
- Harman, G. 1975. Moral Relativism Defended. *Philosophical Review* 84, 3-22.
- 1977. *The Nature of Morality*. Oxford: Oxford University Press.
- 2015. Moral Realism is Moral Relativism. *Philosophical Studies* 172, 855-863.
- Harrison, G.K. 2013. The Moral Supervenience Thesis is Not a Conceptual Truth. *Analysis* 73, 62-68.
- Hempel, C. 1969. Reduction: Ontological and Linguistic Facets. In S. Morgenbesser, P. Suppes, and M. White (ed.), *Philosophy, Science, and Method: Essays in Honour of Ernest Nagel*. New York: St Martin's Press, 179-199.
- Hieronymi, P. 2004. The Force and Fairness of Blame. *Philosophical Perspectives* 62, 529-555.
- Horgan, T. and Timmons, M. 1991. New Wave Moral Realism Meets Moral Twin Earth. *Journal of Philosophical Research* 16, 447-465.
- 1992a. Troubles for New Wave Moral Semantics: The Open Question Argument Revived. *Philosophical Papers* 21, 153, 175.

- 1992b. Troubles on Moral Twin Earth: Moral Queerness Revived. *Synthese* 92, 221-260.
- Hernandez, J.G. 2011. *The New Intuitionism*. London: Continuum.
- Hills, A. 2009. Supervenience and Moral Realism. In A. Heike and H. Gottleib (eds.), *Reduction, Abstraction, Analysis*. Lancaster: Ontos Verlag.
- Hinckfuss, I. 1987. *The Moral Society: Its Structure and Effects*. Discussion Papers in Environmental Philosophy. Canberra: Australian National University.
- Hume, D. 1739-40. *A Treatise of Human Nature*. Oxford: Clarendon Press, 1975.
- 1748. *An Enquiry Concerning Human Understanding*. Oxford: Oxford University Press, 2007.
- 1751. *An Enquiry Concerning the Principles of Morals*. Oxford: Clarendon Press, 1975.
- Huemer, M. 2005. *Ethical Intuitionism*. Hampshire: Palgrave MacMillan.
- Hursthouse, R. 1999. *On Virtue Ethics*. Oxford: Oxford University Press.
- Ingram, S. 2015a. After Moral Error Theory, After Moral Realism. *The Southern Journal of Philosophy* 53, 227-248.
- 2015b. The Moral Fixed Points: Reply to Cuneo and Shafer-Landau. *Journal of Ethics and Social Philosophy* (<http://www.jesp.org>), February 2015.
- Forthcoming. I Can't Relax! You're Driving Me Quasi! *Pacific Philosophical Quarterly*.
- MS. Epistemology Shmeπισtemology: Moral Error Theory and the Problem of Epistemic Normativity.
- Jackson, F. and Pettit, P. 1995. Moral Functionalism and Moral Motivation. *Philosophical Quarterly* 45, 20-40.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- James, W. 1896. The Will to Believe. In *The Will to Believe and Other Essays in Popular Philosophy*. London: Harvard University Press, 1979.

- 1907. *Pragmatism: A New Name for Some Old Ways of Thinking*. Cambridge, MA: Harvard University Press, 1975.
- Joyce, R. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- 2005. Moral Fictionalism. In M.E. Kalderon (ed.), *Fictionalism in Metaphysics*. Oxford: Oxford University Press, 287-313.
- 2006. *The Evolution of Morality*. Cambridge, MA: The MIT Press.
- Kahane, G. 2013. Must Metaethical Realism Make a Semantic Claim? *Journal of Moral Philosophy* 10, 148-178.
- Kant, I. 1785. *Groundwork of the Metaphysics of Morals*. Cambridge: Cambridge University Press, 1997.
- Kelly, T. 2006. The Epistemic Significance of Disagreement. In J. Hawthorne and T. Gendler Szabo (eds.), *Oxford Studies in Epistemology: Volume 1*. Oxford: Oxford University Press, 167-196.
- Kidd, I.J. 2013. Is Naturalism Bleak? *Environmental Values* 22, 689-702.
- Forthcoming. Intellectual Humility, Confidence, and Argumentation. *Topoi*.
- Kim, J. 1993. Concepts of Supervenience. In J. Kim (ed.), *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press, 53-78.
- Kitcher, P. 2011. *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Korsgaard, C.M. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Kramer, M.H. 2009. *Moral Realism as a Moral Doctrine*. Chichester: Wiley-Blackwell.
- Kripke, S. 1980. *Naming and Necessity*. Oxford: Blackwell.
- Laurence, S., Margolis, E. and Dawson, A. 1999. Moral Realism and Twin Earth. *Facta Philosophica* 1, 135-165.

- Lenman, J. 1999. Michael Smith and the Daleks: Reason, Morality, and Contingency. *Utilitas* 11, 164-177.
- 2009. Naturalism Without Tears. In J. Suikkanen (ed.) *Essays on Derek Parfit's On What Matters*. Chichester: Wiley-Blackwell, 21-38.
- 2010. Humean Constructivism in Moral Theory. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 5*. Oxford: Oxford University Press, 175-193.
- 2013. Ethics Without Errors. *Ratio* 26, 391-409.
- Lewis, D. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61, 343-377.
- 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- 1989. Dispositional Theories of Value. *Proceedings of the Aristotelian Society*, Supplementary Volume, 113-137.
- Little, M. Moral Generalities Revisited. In B. Hooker and M. Little (eds.), *Moral Particularism*. Oxford: Oxford University Press, 276-304.
- Lowe, E.J. 1998. *The Possibility of Metaphysics: Substance, Identity and Time*. Oxford: Oxford University Press.
- 2000. *A Survey of Metaphysics*. Oxford: Oxford University Press.
- 2012. What is the Source of Our Knowledge of Modal Truths? *Mind* 122, 919-950.
- MacFarquhar, L. 2011. How to Be Good, Interview with Derek Parfit. *The New Yorker*. Accessed online, Friday 14th November 2014. [<http://www.newyorker.com/magazine/2011/09/05/how-to-be-good>]
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. London: Penguin.
- Majors, B. 2005. Moral Discourse and Descriptive Properties. *Philosophical Quarterly* 55, 475-494.
- Marks, J. 2013. *Ethics Without Morality: A Defence of Amoralism*. London: Routledge.
- McGrath, S. 2008. Moral Disagreement and Moral Expertise. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 3*. Oxford: Oxford University Press, 87-107.

- 2010. Moral Realism Without Convergence. *Philosophical Topics* 38, 58-90.
- 2014. Relax? Don't Do It! Why Moral Realism Won't Come Cheap. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 9*. Oxford: Oxford University Press, 186-214.
- McKeever, S. and Ridge, M. 2006. *Principled Ethics: Generalism as a Regulative Ideal*. Oxford: Oxford University Press
- Miller, C. 2009. The Conditions of Realism. *The Journal of Philosophical Research* 34, 123-155.
- McLaughlin, B. and Bennett, K. 2014. Supervenience. In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
- Moore, G.E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- McPherson, T. 2008. Metaethics and the Autonomy of Morality. *Philosopher's Imprint* 8, 1-16.
- 2011. Against Quietist Normative Realism. *Philosophical Studies* 154, 223-240.
- 2012. Ethical Non-Naturalism and the Metaphysics of Supervenience. In R. Shafer-Landau (ed.) *Oxford Studies in Metaethics: Volume 7*. Oxford: Oxford University Press, 205-234.
- Nagel, T. 1997. *The Last Word*. New York: Oxford University Press.
- Nietzsche, F. 1886. *Beyond Good and Evil*. Oxford: Oxford University Press, 1998.
- 1887. *On the Genealogy of Morals*. Oxford: Oxford University Press, 1996.
- Nolan, D., Restall, G., and West, C. 2005. Moral Fictionalism Versus the Rest. *Australasian Journal of Philosophy* 83, 307-330.
- Nowell-Smith, P.H. 1954. *Ethics*. London: Penguin.
- Nussbaum, M. 1992. Human Functioning and Social Justice: In Defense of Aristotelian Essentialism. *Political Theory* 20, 202-246.
- 2000. Why Practice Needs Ethical Theory: Particularism, Principle, and Bad Behaviour. In B. Hooker and M. Little (eds.), *Moral Particularism*. Oxford: Oxford University Press, 227-255.

- Oddie, G. 2005. *Value, Reality, and Desire*. Oxford: Oxford University Press.
- Olson, J. 2014. *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- 2006. Normativity. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 1*. Oxford: Oxford University Press, 325-380.
- 2011. *On What Matters: Volume 1 and Volume 2*. Oxford: Oxford University Press.
- Plantinga, A. 1974. *The Nature of Necessity*. Oxford: Oxford University Press.
- 2010. Naturalism, Theism, Obligation, and Supervenience. *Faith and Philosophy* 27, 247-272.
- Popper, K. 1959. *The Logic of Scientific Discovery*. London, Routledge.
- Quine, W.V.O. 1951. Two Dogmas of Empiricism. *Philosophical Review* 60, 20-43.
- Ratcliffe, M. 2005. The Feeling of Being. *Journal of Consciousness Studies* 12, 43-60.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Ridge, M. 2007. Anti-Reductionism and Supervenience. *Journal of Moral Philosophy* 4, 330-348.
- 2014. *Impassioned Belief*. Oxford: Oxford University Press.
- 2015. I Might Be Fundamentally Mistaken. *Journal of Ethics and Social Philosophy* (<http://www.jesp.org>), October 2015.
- Roberts, R.C. and Wood, W.J. 2006. *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Oxford University Press.
- Rowland, R. 2013. Moral Error Theory and the Argument from Epistemic Reasons. *Journal of Ethics and Social Philosophy* 7, 1-24.
- Rubin, M. 2015. Normatively Enriched Moral Meta-Semantics. *Philosophy and Phenomenological Research* 91, 386-410.
- Ruse, M. 1986. *Taking Darwin Seriously*. Oxford: Blackwell.

- Russell, L. 2006. Evil-Revivalism Versus Evil-Skepticism. *The Journal of Value Inquiry* 40, 89-105.
- Scanlon, T.M. 2008. *Moral Dimensions: Permissibility, Meaning, and Blame*. Cambridge, MA: Harvard University Press.
- 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schaffer, J. 2004. Two Conceptions of Sparse Properties. *Pacific Philosophical Quarterly* 85, 92-102.
- Schroeder, M. 2005. Realism and Reduction: The Quest for Robustness. *Philosopher's Imprint* 5, 1-18.
- 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- 2014. The Price of Supervenience. In *Explaining the Reasons We Share: Explanation and Expression in Ethics*. Oxford: Oxford University Press, 124-144.
- Shafer-Landau, R. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- 2004. *Whatever Happened to Good and Evil?* Oxford: Oxford University Press.
- 2006. Ethics as Philosophy: A Defence of Ethical Non-Naturalism. In T. Horgan and M. Timmons (eds.), *Metaethics After Moore*. Oxford: Oxford University Press, 209-232.
- 2009. A Defence of Categorical Reasons. *Proceedings of the Aristotelian Society* 109, 189-206.
- 2012. Evolutionary Debunking, Moral Realism and Moral Knowledge. *Journal of Ethics and Social Philosophy*, www.jesp.org, 7, 1-37.
- Shemmer, Y. 2007. Desires as Reasons. *Philosophy and Phenomenological Research* 75, 326-348.
- Sidgwick, H. 1907. *The Methods of Ethics*. London: Hackett, 1981.
- Skarsaune, K.O. 2011. Darwin and Moral Realism: Survival of the Iffiest. *Philosophical Studies* 152, 229-243.

- 2015. How To Be a Moral Platonist. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 10*. Oxford: Oxford University Press, 245-272.
- Smith, M. 1994. *The Moral Problem*. Oxford: Blackwell.
- Sobel, D. 2009. Subjectivism and Idealization. *Ethics* 119, 336-352.
- Sripada, C. and Stich, S. 2006. A Framework for the Psychology of Norms. In Carruthers, P., Laurence, S., and Stich, S. (eds.), *The Innate Mind, Volume 2: Culture and Evolution*. Oxford: Oxford University Press.
- Stevenson, C. 1944. *Ethics and Language*. New Haven, CT: Yale University Press.
- Stoljar, D. 2010. *Physicalism*. New York: Routledge.
- Stratton-Lake, P. 2002. Introduction. In W.D. Ross, *The Right and the Good*. Oxford: Oxford University Press.
- Street, S. 2006. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies* 127, 109-166.
- 2008. Constructivism about Reasons. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 3*. Oxford: Oxford University Press, 207-245.
- 2009. In Defence of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues* 19, 273-298.
- 2010. What is Constructivism in Ethics and Metaethics? *Philosophy Compass* 5, 363-384.
- 2011. Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 6*. Oxford: Oxford University Press, 1-32.
- 2012. Coming to Terms with Contingency: Humean Constructivism about Practical Reason. In J. Lenman and Y. Shemmer (eds.), *Constructivism in Practical Philosophy*. Oxford: Oxford University Press, 40-59.
- 2015. Does Anything Really Matter or Did We Just Evolve to Think So? In G. Rosen, A. Byrne, J. Cohen, and S. Shiffrin (eds.), *The Norton Introduction to Philosophy*. New York: Norton, 685-693.

- Streumer, B. 2008. Are There Irreducibly Normative Properties? *Australasian Journal of Philosophy* 86, 537-561.
- 2011. Why there really are no irreducibly normative properties. In D. Bakhurst, B. Hooker and M. Little (eds.), *Thinking about Reasons: Themes from the Philosophy of Jonathan Dancy*. Oxford: Oxford University Press, 310-336.
- Sturgeon, N. 1988. Moral Explanations. In G. Sayre-McCord (ed.), *Essays on Moral Realism*. Cornell: Cornell University Press, 229-255.
- 2006. Ethical Naturalism. In D. Copp (ed.), *The Oxford Handbook of Ethical Theory*. Oxford: Oxford University Press, 91-121.
- 2009. Doubts about the Supervenience of the Evaluative. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 3*. Oxford: Oxford University Press: 53-90.
- Suikkanen, J. 2010. Non-Naturalism: The Jackson Challenge. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 5*. Oxford: Oxford University Press, 53-90.
- Tännsjö, T. 2010. *From Reasons to Norms: On the Basic Question in Ethics*. Dordrecht: Springer.
- Thomson, J.J. 2001. *Goodness and Advice*. Princeton: Princeton University Press.
- 2008. *Normativity*. Chicago: Open Court.
- Tiberius, V. 2012. Open-Mindedness and Normative Contingency. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 7*. Oxford: Oxford University Press, 182-204.
- Tropman, E. 2008. Naturalism and the New Moral Intuitionism. *Journal of Philosophical Research* 33, 163-184.
- 2009. Renewing Moral Intuitionism. *Journal of Moral Philosophy* 6, 440-463.
- Unwin, N. 1999. Quasi-Realism, Negation, and the Frege-Geach Problem. *The Philosophical Quarterly* 49, 337-352.
- van Roojen, M. 2006. Knowing Enough to Disagree: A New Response to the Moral Twin Earth Argument. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 1*. Oxford: Oxford University Press, 161-193.

- Vavova, K. 2014. Debunking Evolutionary Debunking. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics: Volume 9*. Oxford: Oxford University Press, 76-101.
- 2015. Evolutionary Debunking of Moral Realism. *Philosophy Compass* 10, 104-116.
- Väyrynen, P. 2013. Grounding and Normative Explanation. *Proceedings of the Aristotelian Society*, Supplementary Volume, 87, 155-178.
- Wallace, R.J. 2011. Dispassionate Opprobrium: On Blame and the Reactive Sentiments. In R.J. Wallace, R. Kumar, and S. Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: Oxford University Press, 348-372.
- Watson, G. 1996. Two Faces of Responsibility. *Philosophical Topics* 24, 227-248.
- Wedgwood, R. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- Wielenberg, E.J. 2010. On the Evolutionary Debunking of Morality. *Ethics* 120, 441-464.
- 2011. The Failure of Brown's Supervenience Argument. *Journal of Ethics and Social Philosophy* (<http://www.jesp.org>), October 2011.
- 2014. *Robust Ethics: The Metaphysics and Epistemology of Godless Normative Realism*. Oxford: Oxford University Press.
- Williams, B. 1981a. Practical Necessity. In *Moral Luck*, Cambridge: Cambridge University Press, 124-131.
- 1981b. Internal and External Reasons. In *Moral Luck*, Cambridge: Cambridge University Press, 101-113.
- 1985. *Ethics and the Limits of Philosophy*. London: Routledge.
- Williamson, T. 2001. Ethics, Supervenience and Ramsey Sentences. *Philosophy and Phenomenological Research* 62, 625-630.
- Wilson, J. 2010. What is Hume's Dictum, and Why Believe It? *Philosophy and Phenomenological Research* 80, 595-637.
- 2014. Hume's Dictum and the Asymmetry of Counterfactual Dependence. In A. Wilson (ed.), *Chance and Temporal Asymmetry*. Oxford: Oxford University Press, 258-279.

- 2015. Hume's Dictum and Metaphysical Modality: Lewis's Combinatorialism. In B. Loewer and J. Schaffer (eds.), *The Blackwell Companion to David Lewis*. Oxford: Blackwell, 138-158.
- Wolf, S. 2011. Blame, Italian Style. In R.J. Wallace, R. Kumar, and S. Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: Oxford University Press, 332-347.
- Wood, A. 2011. Humanity as an End in Itself. In D. Parfit, *On What Matters: Volume 2*. Oxford: Oxford University Press, 58-81.
- Woods, J. 2014. Expressivism and Moore's Paradox. *Philosophers Imprint* 14, 1-12.
- Wright, C. 1992. *Truth and Objectivity*. London: Harvard University Press.