Normative Foundations

Kevin Austin Zhang

Dissertation submitted for the DPhil in Philosophy at the University of Oxford

Friday, July 15, 2022 (Trinity W12)

Abstract

This dissertation presents a theory of normativity. The central claim to be advanced is that human beings are constitutively committed to certain objective moral and other normative requirements. Furthermore, setting itself against the proposition that there are brute and mind-independent normative facts, the dissertation also advances positive theses about how normativity arises from the mind, as well as why a mind-given theory of normativity need not be nihilistic. After advancing this positive vision of normativity, the dissertation turns to examine prominent arguments for moral objectivity from Christine Korsgaard, Michael Smith, Julia Markovits, and Sharon Street.

The first half of the dissertation, from Chapters 1 through 5, develops the positive theory of normativity. The second half of the dissertation, from Chapters 6 through 10, engages with existing arguments for moral objectivity from the recent literature. The result is the presentation of an interconnected set of views on certain basic questions about normativity, as well as a constructive and critical examination of recent meta-ethical work.

Contents

1 Fundamentals

1.1 Introduction

1.2 Realism and Its Discontents

1.3 Normativity from the Mind

1.4 Notes on Existing Literature

1.5 The Project Ahead

2 The Normative World

2.1 Introduction

2.2 Against Attitude Monism

2.3 The Positive Picture

2.4 Theoretical Upshots

2.5 Normative Lens

3 The Structure of Normativity

3.1 Introduction

3.2 The Standard of Possibility

3.3 Endorsement and Representation

3.4 Substantive Upshots

4 Nihilistic Concerns

4.1 Introduction

4.2 Value and the Import of Potentialities

4.3 Mind-Given Irreducibility

4.4 The Reality of Normativity

5 Moral Objectivity

5.1 Introduction

5.2 The Conceptual Argument

5.3 The Empirical Argument

5.4 The Demands of Love

6 Korsgaard's *The Sources of Normativity* (1996)

6.1 Introduction

6.2 Valuing Humanity: Korsgaard's Argument

6.3 Valuing Humanity: Street and Bukoski

6.4 The Problem: A Non-Sequitur

7 Korsgaard's *Self-Constitution* (2009)

7.1 Introduction

7.2 The Publicity of Reasons: Korsgaard's Argument from *Self-Constitution*

7.3 The Publicity of Reasons: Korsgaard's Argument from *Sources*

8 Smith's Constitutivism (2013)

8.1 Introduction

8.2 Smith's Argument

8.3 The Desires of the Ideal Agent

8.4 The Dispositional Claim

9 Markovits' *Moral Reason* (2014)

9.1 Introduction

9.2 Markovits' Argument

9.3 Understandability to Others

9.4 Lurking Realism

10 Street's Buddhist Approach (2016)

10.1 Introduction

10.2 Street's Proposal

10.3 Three Comments

11 Conclusion

Acknowledgements

This dissertation is the product of a significant amount of generosity. I would first like to thank my primary supervisor, Professor Roger Crisp, for being the best supervisor that any graduate student could ask for. He has not only provided insightful and sharp comments on every chapter of this dissertation, but he has also been a philosophical model for his charity and interest towards the development of different and burgeoning views. My constant feeling of being supported and encouraged speaks to a fundamental kindness and wisdom in his character – and I remain deeply indebted to him, for his steadfast generosity towards me, and for encouraging me from the early onset of this project, up until its very final stages.

I would next like to thank my second supervisor, Professor Ruth Chang, for graciously signing on to supervise this project amidst an incredibly demanding schedule: I have been most grateful to her for her insightful comments, advice, and encouragement on the dissertation, all of which were given in a supererogatory and generous way.

I must then thank a number of warm and supportive friends. For those who read very substantial amounts of the dissertation and provided insightful comments and suggestions, I am grateful to John Fan, Avital Fried, Will Siegmund, Denis Kazankov, Leo Schneider, Angelo Ryu, Jyoti Nayak, Crystal Lam, Yoni Genende, and Konstantinos Konstantinou. I have moreover been blessed to receive help and insight from a number of generous individuals on specific segments of the dissertation: Edison Yi, Sarah Hirschfield, Lily Moore-Eissenberg, Alexander Lam, Clemente Recabarren, Fernando Contreras, Albert Zuo, Nina Wang, Shiye Su, Emerson Salovaara, and Vicky Yu. Special thanks goes to Cayla Bleoaja for her helpful reading of the dissertation, as well as for her enthusiastic and sustained encouragement.

There are three more individuals who I would like to thank for their influence on this work. The first two are Christine Korsgaard and Sharon Street – two philosophers who had a seminal influence on my encounter with meta-

ethics, and to whom I owe a grateful intellectual debt. The third person is Amanda Cheng, whose intellectual and personal influence on the development of this dissertation has, over the years, had a significant impact on both the project and its author.

It would be regrettable if I did not also recall and thank a handful of generous individuals who have helped to encourage me towards Oxford in the first place, and who have made this time a wonderful one for me: Ruby Shao and David Oakley were two friends who encouraged and helped me early on, and I also want to thank my friends at the Canterbury Institute for their support in many ways, with special thanks to Dominic Burbidge for his quiet kindness and wisdom.

Finally, I wish to thank my family for their unconditional support and love: Mom, Dad, Albert, Brandon, Mimi, and Lao Lao. Growing up, I think that none of us would have expected me to find myself writing a doctorate in philosophy one day, and I am happy to report that this surprise has been a welcome one for all of us.

1 Fundamentals

1.1 Introduction

Normativity is one of the defining features of human life. In fact, apart from consciousness and existence itself, normativity is arguably what most profoundly characterizes what it is like to be a human. This is because normativity encompasses all that gives existence a *point*, from goodness to beauty to meaning to morality.

The fact that philosophers now employ the concept of normativity to encompass these significant and disparate notions is itself something worth wondering about. After all, from a historical and linguistic standpoint, the word 'normative' has been neither timeless nor widespread. In the English language, its first recorded usage dates back only to the mid to late 1800s, and the word did not become more prominent until the latter half of the 20th century.[1] In other languages, there is not necessarily a precise correlate for the term. So what justifies the recent understanding that notions like goodness, beauty, meaning, and morality can all be understood under the umbrella of normativity?[2]

What seems to unite these different notions is the fact that all of them have to do with giving the world some *directionality* to it. As I have said, normativity encompasses everything that gives our lives a *point* – it is what makes it the case that we can orient ourselves in the directions of good and bad, understand that there are things that we *should* do, and find some of our experiences to be beautiful and valuable and meaningful. This does not amount to a precise definition of normativity, for I suspect that it is what philosophers call a basic concept – one with a distinct character that can be elucidated through

[1] The Oxford English Dictionary and Merriam-Webster have the first known use of 'normative' at 1852, whereas Dictionary.com has it at 1875-1880. 'Normativity' only appears for the first time, according to the Oxford English Dictionary, in 1935. Its recent increase in usage is based on historical data from Google Books Ngram Viewer.

[2] Throughout this dissertation, I therefore use the term 'normativity' in a broad sense. For some examples of similar usage in the recent literature, see Sharon Street (2008: 209n5) and T. M. Scanlon (2021: 25), among others.

analogy or explication, but which cannot be defined in a way that does not simply point back to itself. My preferred way of illustrating the concept of normativity is to imagine a world in which there are no sentient creatures in it, no God, and no other facts about goodness, beauty, or rightness; there are only atoms, rocks, and material objects that move according to causal laws. This is a world that has no normativity in it – it is a world that gives the fullest sense to Nietzsche's fear of there being no "up or down."[3] There is some fortune to the fact that we live in a normative world, for our world contains at least the promise of goodness, beauty, meaning, and morality.

But once we reflect upon normativity, we encounter certain perennial philosophical questions: where does normativity come from, and are there any *objective truths* about normativity? Most of us, apart from the skeptical few, think that there are objective truths about there being physical objects that exist outside of our minds, or that there are certain scientific laws that govern the world. But the question at hand concerns an analogue in the normative domain, and it confronts us through both practical and theoretical life: are there objective truths about what is good, what is beautiful, what is meaningful, and what is morally right?

Relativism about normativity is likely as old as philosophy itself, and we can count everyone from Thrasymachus to Mackie as those who have doubted the objectivity of normativity. But even so, I take it that most people have at least *some* pre-theoretic impulse to believe that there are objective normative truths: that there are certain truths about normativity that are independent of the contingent thoughts or opinions that people come to adopt.[4] Nowhere is this more vivid than in the case of morality. Many of my readers, upon witnessing the cruel torture of an infant, would feel at least some instinct to believe that,

[3] See Nietzsche in *The Gay Science* (1882: Section 125).

[4] We will therefore be interested in *objectivity* in the sense that someone can make a *mistake* about some normative matter, by the lights of some standard that has a significant degree of independence from the particular thoughts or other attitudes of that individual and other persons. This can be contrasted with a different sense of objectivity that concerns complete metaphysical independence from the mind. For more on the different senses of objectivity, see §1.3, including ft. 24.

even if most of the world came to commend this action as *right*, it would still be the case that this action was *wrong.* Or, think about the matter this way: imagine that Hitler had won the war, and a strong majority of the world's population came to believe that his specific ideology of genocide was correct. Most of my readers will feel at least some pre-theoretic impulse to maintain that, even in such a world of collective assent to Hitler, his actions would never have been right – they would still have been profoundly *wrong.* And this would mean that there must be certain truths about normativity which obtain independently of the particular judgments that individuals or collectives could come to adopt.

It is because of this pre-theoretic impulse towards objectivity that relativism will never prove to be an ultimately satisfying position: no matter how intellectually coherent the view might seem, or how strong the objections to objectivity may appear to be, relativism will always leave our deeper philosophical aspirations in a contorted place. That is because there is a more profound yearning in us to vindicate that there are indeed certain objective normative truths. It is perhaps for this reason that even J. L. Mackie, one of the leading relativists of the recent decades, himself acknowledged that "the denial of objective values can carry with it an extreme emotional reaction, a feeling that nothing matters at all, that life has lost its purpose" (1977: 34).

Historically, the two most obvious paths to defending normative objectivity have been to appeal to God, or to assert that there are certain normative facts, properties, or entities that exist independently of the human mind – facts, properties, or entities which do not *arise from* any of the activities or attitudes of the human mind.[5] We can call the first *theism*, and the second *realism*. Of course, these positions are not incompatible: one could be a theistic realist, maintaining that God created the mind-independent normative facts, properties, or entities. But importantly, these positions are also divorceable: one could believe that God directly authors normative requirements, without

[5] It would also be possible to say that these normative facts, properties, or entities are not *grounded* in the mind, but because the notion of grounding is a contested one, I shall often prefer to avoid it.

creating any normative facts, properties, or entities that are metaphysically distinct from the divine will; or one could refrain from believing in God, as many secular realists do, but nevertheless maintain that there are these mind-independent normative facts, properties, or entities.

This dissertation does not claim that either theism or realism is false. These are both extremely important positions, with much to be said for them, and I am also skeptical that they could ever be disproven. Rather, I have a different and softer aim: I want to ask the question of whether normative objectivity can be vindicated *without* appealing to either theism or realism. And I believe that there are both intrinsic and practical reasons that motivate such a project.

Theism continues to be the reason why much of the world believes in normative objectivity. However, even so, it has long remained a compelling question whether normative objectivity can be vindicated without appealing to God. For instance, Aquinas asked this question, and he concluded that normative requirements – embodied in natural law – could be shown to be objective through natural reason alone. Similarly, Grotius famously asked whether we could show that there are moral requirements, "even if we should concede that which cannot be conceded without the utmost wickedness, that there is no God" – and his conclusion was that we could, at least to some extent.[6] So, with Aquinas and Grotius, even a theist can be interested in the question of whether normative objectivity can be defended, prior to considerations about God. But apart from the intrinsically compelling nature of the question, there is also a practical reason for our inquiry: much of the modern world's population *is* secular, yet these people still want to know, and need to settle, whether morality and other normative matters are objective. Absent some deeper theological resolution in our world, many people are wondering whether philosophical reflection can show that there are objective normative truths.

[6] See Grotius (1625 [2003]: 92), in Schneewind.

Something similar could be said about realism – it is an intrinsically interesting question whether normative objectivity can be vindicated without appealing to mind-independent normative facts, and this is also of practical import, because at least some people find it hard to believe in the existence of such mind-independent normative facts. But with realism, there is some more that I want to say.

1.2 Realism and Its Discontents

Realism, I have said, is the position that there are certain normative facts, properties, or entities that exist independently of the human mind. This definition is broad, because realism has been championed in numerous ways since its early paradigmatic expression in Plato, for whom these mind-independent normative entities were the Forms. In more recent centuries, realism has been defended by Samuel Clarke and Richard Price in the 18th century, Henry Sidgwick in the 19th century, G. E. Moore, W. D. Ross, and H. A. Prichard in the 20th century, and Derek Parfit and T. M. Scanlon in the 21st century, among many others.[7] Because realism is one of the most obvious routes towards defending normative objectivity, I suspect that it will continue to resurface in every new philosophical generation.

In seeking to understand realism, it can be illustrative to imagine how a case of normative disagreement comes to an end for a realist. Suppose that you are having a fundamental disagreement with someone about what is good, beautiful, or right. You both establish that there is no disagreement over the logical and empirical facts, that you have both meditated upon the issue for a long time, and that the normative disagreement cannot be explained by way of some more *basic* normative disagreement: there simply remains a *pure* and *fundamental* normative disagreement. Let us also assume that this is a case

[7] Other recent defenses of realism can be found in Murdoch (1970), Nagel (1986), Dancy (2000), Shafer-Landau (2003), Crisp (2006), Skorupski (2010), Enoch (2011), and Dworkin (2011), among other works.

where the realist believes there is a mind-independent fact of the matter. According to the realist, what is happening here is that one of you is correctly intuiting or accessing the mind-independent normative truth, whereas the other is somehow failing to. If you believe that *you* are the one who has successfully hit upon that mind-independent normative truth, you might ask your interlocutor to explain herself; and it will seem that the best she can say is something like the following: "It just seems like this *is* the normative truth – it couldn't be more obvious to me. Don't you see the same?"

Although this is only an illustration, it is this demand for ultimate explanation that offers an important insight into understanding the realist: at least for the realist who does not further appeal to God, the asserted normative truth will turn out to be a *brute* fact. By this, I mean that most realists must insist that the existence of the mind-independent normative truth is not something that is open to further explanation. For traditional realists like Plato, who envisioned that there were normative properties or entities in some abstract realm, it would be the metaphysical existence of these properties or entities that would not be further explained; for more recent realists like Parfit and Scanlon, who have attempted to de-emphasize the metaphysical commitments involved in their accounts, it would still be that the fundamental normative facts are themselves not open to further explanation. So, for instance, in capturing the bruteness involved in Scanlon's account, Gideon Rosen writes that "the basic normative principles are brute synthetic facts *with no explanation*" (2017: 862-863, emphasis added); and Ruth Chang, in characterizing how the realist conceives of certain normative reasons, writes that "there is no more explanation to be had… the reason is ungrounded" (2021: 97-98). Parfit, describing the bruteness involved in his own view, speaks of the "irreducibility" of normative truths (2011b: 267).[8]

Notice that the two notions I have used to characterize realism thus far, those of *mind-independence* and *bruteness*, have a close relationship with one another. If some normative truth is *mind-independent*, then it cannot be further

[8] See also Scanlon (2014: 210).

explained as arising from the mind; if some normative truth is *brute*, then it cannot be further explained at all. This means the two notions can often travel together. For instance, *if* some normative truth is brute, then it will *necessarily* also be mind-independent: for if it cannot be further explained at all, then it will certainly not be further explained with reference to the mind. Alternatively, *if* some normative truth is mind-independent, then it will *likely* also be brute: for if it cannot be further explained through reference to the mind, then it is likely that the theorist will maintain that it has no further explanation at all. But there is at least one obvious case in which this final inference will not follow, and that is if the realist further appeals to God to explain the existence of the mind-independent normative truth. Because I shall mostly be engaging with secular realists in the pages to follow, from here onwards I will use the term 'realism' to primarily refer to these secular authors, for whom the mind-independent normative truths are also *brute* ones – though the reader should take note of the important caveat made here.

Because realism short-circuits explanation in this way, it can lead to a number of discontents. First, the bruteness that is involved in realism means that it is *dogmatic* in some important sense. Realism is dogmatic in the sense that, when it tries to establish the existence of some basic normative truth, it cannot ultimately do more than simply *insist* upon the truth of that normative proposition, without offering some deeper explanation for the existence of that truth. While this observation does not show that realism is *false*, it does capture an important reason why realism can seem unhelpful. Here is Christine Korsgaard:

> If someone falls into doubt about whether obligations really exist, it doesn't help to say 'ah, but indeed they do. They are *real* things'. Just now he doesn't see it, and herein lies his problem. To see this, go back to the case where you are being asked to face death rather than do a certain action. You ask the normative question: you want to know whether this terrible claim on you is justified. Is it really true that this is what you *must* do? The realist's answer to this question is simply 'Yes'. That is, *all* he can say is that it is *true* that this is

> what you ought to do. This is of course especially troublesome when the rightness of the action is supposed to be self-evident and known through intuition, so that there is nothing more to say about it… And that is the problem with realism: it refuses to answer the normative question. It is a way of saying that it cannot be done. (Korsgaard 1996: 38, 40)

Importantly, many realists can be read as maintaining that there are mind-independent normative truths that apply to an individual, even if that individual cannot ever come to recognize those normative truths through a process by which we reason with him.[9] In this way, realism posits the existence of mind-independent normative truths that exert a binding force upon how we ought to live, even though we could fail to ever come to see the validity of those normative truths from our perspectives.[10]

Realism can also lead to metaphysical and epistemological concerns. Recall that realism posits normative truths that are independent of the human mind. *Metaphysically* then, we can wonder what these mind-independent normative truths *are*; and *epistemologically*, we can wonder what sort of faculty we would need to access these mind-independent truths. Here is how Mackie memorably put these points:

> [The realist's values] would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else. (Mackie 1977: 38)

No doubt there have been replies to Mackie, but I believe his concerns still hold some intuitive force, especially towards theorists who hold secular worldviews

[9] As Russ Shafer-Landau puts it: "Realists believe… that the moral standards that fix the moral facts are not made true by virtue of their ratification from within any given *actual* or *hypothetical* perspective" (2003: 15, emphasis edited). See also Parfit (2011b: 271).
[10] For an objection along these lines, see also Korsgaard (2009: 6), Markovits (2014: 56 & 136, 2016: 526), and Markovits & Walden (2021: 319).

that otherwise emphasize naturalistic explanations.[11] In addition to these traditional concerns, there is another problem: if normative truths are entirely independent of the human mind, we can wonder why we should expect a general *alignment* between these mind-independent normative truths and our actual basic normative commitments. Unless the mind-independent normative truths exert a direct causal force upon the formation of our basic normative commitments, those basic normative commitments will have been largely shaped by a set of orthogonal influences, such as evolution or cultural formation. But then the question emerges: if these *orthogonal* causes are what led us to have the basic normative commitments that we have, then how can we ensure that our basic normative commitments are not significantly off-track from the *mind-independent* domain of normative truths? This is how I understand the thrust of Street's influential evolutionary debunking argument (2006), which I defend at length elsewhere.[12]

These constitute some of the standard objections to realism.[13] But there is an even more basic question one can ask the realist: where do the mind-independent normative truths *come from*? Remember that these are facts of the most important kind, which govern how we are supposed to live, and which concern everything that *matters* to us. So do these normative facts have some kind of origin? Of course, the realist may insist that this question is confused: after all, there are many sorts of facts, like mathematical or logical or scientific facts, that have to bottom out in bedrock somewhere, where there is not much room for entertaining a further request for explanation. Even so, it has always seemed to me that normative truths are in at least some respects different from mathematical, logical, and scientific facts: normativity governs the specific ways

[11] For some responses to Mackie, see Brink (1984) and Shepski (2008); see Garner (1990) for a response to Brink.
[12] See *On Moral Objectivity* (Zhang 2019: Chapter 2).
[13] There are other well-known objections that have to do with *motivation* and *disagreement* (see Williams 1995: 39 and Mackie 1977: 36, respectively). I will not discuss these objections here, at least in part because I do not find them as persuasive. For some skeptical discussion of the motivational objection, in particular, see Markovits (2014: Ch 2), Sobel (2016: Ch 7), and Hawkins (2019: 96). For helpful summaries of objections to realism in general, see Kolodny (2005: 511), Markovits & Walden (2021: 318-320), and Scanlon (2021: 25-28).

in which our form of life – which has only existed for a small fraction of the time that the Universe has been around – is supposed to live; whereas mathematical, logical, and scientific facts seem to have *some* more impersonal connection to how the Universe has always functioned.[14] Is it not odd to think that there have always existed these mind-independent normative truths, which exert a binding force upon how we creatures should conduct their lives, yet there is *no* explanation at all for where these truths come from? The refusal to provide such an explanation creates a certain kind of discontent – and it is to be contrasted with John Rawls' suggestion that even an *objective* point of view "must always be *from somewhere*" (1993: 116, emphasis added). On this score, I find the theistic realist somewhat more understandable than the secular realist, for positing God would at least provide some explanation as to why there are mind-independent normative truths in the first place.

Now again, in recounting some of these discontents, I have not claimed that realism is false. There is much that a realist could say in response to these concerns, and I am not confident that these objections, even in their best forms, could ever *disprove* realism.[15] Indeed, one could even take realism to be true on *faith*, which can sound like what some realists are saying at times.[16] I have recounted these discontents in order to show why some theorists may be wary of realism, and to motivate our project of asking whether we might search for a different way of vindicating normative objectivity. After all, note that both

[14] This does not mean that there is no mystery at all to the mathematical, logical, and scientific facts that there are – whether we are children or adults, we can still ask about their origin as well. My point is that there seems something *distinctively* appropriate in asking these questions about normativity, which governs something quite specific in the Universe – our lives. See McGrath (2014) for discussion of some of the disanalogies between normativity and mathematics.

[15] For instance, compare how some critics of realism have claimed that realism is *false*. Street writes that "realism about normative reasons is *false*" (2016a: 165, emphasis edited), and that "there are no mind-independent normative truths" (2016a: 173); see also Street (2008: 207). Williams has also written that "there are only internal reasons for action" (1995: 35), with the implication therefore being that there are *no* external reasons for action.

[16] This might sound like a strong claim, but consider the following quotes. David Enoch, in responding to Street's debunking argument, writes that "I believe that a *miracle* remains" (Enoch 2011: 173, emphasis added); and Ronald Dworkin, in responding to an objection about disagreement, writes that our own moral judgments "may have *nothing* behind them but the *bare and unsubstantiated* conviction that our capacity for moral judgment functions better than theirs did" (Dworkin 1996: 122, emphasis added).

realists and non-realists can agree that a theory of normativity should be assessed with holistic criteria in mind:

> What we should look for, I submit, is the philosophical theory that is best as a theory overall – and this is consistent, of course, with its losing some plausibility points on this or that issue, as long as it makes up for this loss with the plausibility points it honestly earns on other issues. (Enoch 2011: 14)

> The realism/antirealism debate must be decided on holistic considerations. The mere fact that realism brings with it some apparent problems – even very serious ones – is not enough to defeat the view if the alternative, antirealism, is even more implausible in comparison. (Street 2016b: 326)

If realism has some discontents, one way to proceed would be to try to save it from those objections, which is, for instance, what Scanlon attempts to do in *Being Realistic About Reasons* (2014); but another approach would be to see whether some of the attractions of realism, i.e. objectivity, can be recovered without incurring those same objections in the first place. In searching for objectivity in this way, I find much resonance with David Velleman:

> The point is to show that we *can* account for the objectivity of morality *without* positing a normative reality of which judgments of right and wrong can be true. The point is not to *disprove* the existence of such a reality but merely to show that we don't *need* it in order to explain how morality can be objective. (Velleman 2009: 116, emphasis added)

In other words, we can describe our project as seeking to discover what can be said about normativity, *prior to* our taking on any weightier assumptions about theism or realism. This does not mean that we need to *rule out* that there could be further principles or norms which would come from theism or realism being true – it just means we want to investigate what can be demonstrated without adopting these assumptions yet. Metaphorically then, our project can be

regarded as seeking a *floor* for normativity without theism or realism, rather than advancing some incompatible *alternative* to theism or realism. And that is a project that everyone, theists and realists included, can be interested in.

1.3 Normativity from the Mind

If normativity is not to be understood as something that is independent of the mind, then the natural place to turn is somewhere that we are all already familiar with: to understand normativity as a phenomenon that arises from the first-personal standpoint of being a human person. Starting from a young age, we become acquainted with what it is like to experience things as good, as beautiful, as meaningful, and as valuable – and in these encounters, which concern how our minds come to *relate* to the world in a normative dimension, we can find the building blocks for a fuller account of normativity. This is what I shall call a *mind-given* theory of normativity: one that makes essential reference to the mind's relationship to the world, in its explanation of the origin and nature of normativity.[17]

Normativity, we might say, exists from the standpoint of conscious *being* – from the having of a *point of view* on the world. If we were to tell a story about normativity, then normativity came into the world along with the emergence of such conscious points of view that had a valence to them, from which a creature could relate to the world in a normative manner; and without the existence of such standpoints, there would be no normativity.[18] We can also call this the having of a *practical* standpoint – a conscious standpoint that involves at least some thin layer of normative experience to it. Here is Street on the practical standpoint:

---

[17] One will notice that I often prefer to use the metaphor of *relating* to the world, rather than metaphors such as *projecting* or *expressing* or *conferring* or *bestowing* normativity onto the world. Though I sometimes use these latter metaphors as well, what is appealing about the metaphor of *relating* is that it does not lose essential reference to the fact that there is always some *thing* in the world that we are having some normative reaction to, such that normativity inheres at the relational intersection between *subjects* and *objects* of their normative attention.

[18] At least without the stronger assumptions of realism or theism.

> The practical point of view is the point of view occupied by any creature who takes at least some things in the world to be good or bad, better or worse, required or optional, worthy or worthless, and so on – the standpoint of a being who judges, whether at a reflective or unreflective level, that some things call for, demand, or provide reasons for others. (Street 2010: 366)

What distinguishes a merely *conscious* standpoint from a *practical* standpoint is that the latter involves at least *some* dimension of normative experience, even if that is very rudimentary. For instance, insofar as we could conceive of a conscious standpoint that did not feel even a minimal level of positivity or negativity towards the world – a conscious standpoint with no *valence* to its experience at all – then we would not call such a standpoint a *practical* one in our sense.[19] We could say that it was conscious, but not *normatively* conscious. Of course, many creatures could then have practical standpoints: not only do humans have normative experiences, but presumably elephants, cats, and many other animals do as well.[20] We should also be careful to note that it is conceivable that there could be some creatures that have a normative dimension to their experience, even though they do not have the *concepts* that we have to represent normative experience to themselves. That is to say, in the most rudimentary forms of normative consciousness, there may be such a thing as *non-conceptual normativity*, in which there is a positive or negative valence to

[19] In talking about consciousness, it may also be helpful to stipulate that we have in mind what Ned Block calls *phenomenal* consciousness, rather than *access* consciousness. As he puts the distinction:

> Phenomenal consciousness is experience; the phenomenally conscious aspect of a state is *what it is like* to be in that state. The mark of access-consciousness, by contrast, is availability for use in reasoning and rationally guiding speech and action. (Block 1995: 227, emphasis added)

There is nothing *it is like* to be a rock, and so a rock cannot have a conscious or practical standpoint; to that extent, phenomenal consciousness counts as a necessary condition for having a practical standpoint in our sense. Why are we not interested in access consciousness? Because the fundamental insight into normativity, from our mind-given perspective, is that normativity inheres in there being *something it is like* to be a certain creature.

[20] See also Street (2008: 222n33).

the creature's experience – there is something like pain – even though that creature is not a thinking or cerebral being, who organizes experiences through conceptual schemes. These creatures would still have normative consciousness, and therefore a practical standpoint.

Such a picture is sometimes called an *antirealist* conception of normativity, and it is usually thought to then lead to some form of relativism – for if normativity is a phenomenon that essentially depends on the mind, it can be hard to see how anything about normativity could turn out to be objective. But that is the seemingly paradoxical task that we, along with perhaps Kant and some other recent philosophers, are interested in: to at once explain that there are objective normative truths – which obtain independently of our particular thoughts and other attitudes, and of which we can be seriously mistaken – while also explaining normativity as a phenomenon that is not entirely independent of human valuing and activity.[21] That brings us to our central question: how is such a defense of objectivity possible at all?

The strategy is to look at whether there are some things that we humans *already* do, in bringing normativity into the world, that necessarily commit ourselves to certain deeper normative principles or standards. For instance, I will often be interested in arguments that seek to look at the very nature of valuing – whether that is first directed at caring for some garden, or accomplishing some Mafia aim – and to then show that such valuing, no matter where it starts, conceptually or constitutively commits the individual to some stronger moral or normative principle. These normative principles will be *objective* in the sense that, if one has certain features in one's practical standpoint, like valuing at all, then one *just will* be committed to the normativity of these deeper principles and standards, insofar as our arguments are correct – and certainly, though a person cannot escape being normatively bound to these principles, she can still forget about them, knowingly go against them, or fail to realize that she is indeed

[21] There is some disagreement in the literature as to what Kant's meta-ethical position was – in a recent manuscript, I argue that he was neither a standard realist nor a Kantian constructivist (Zhang, "On Kant's Meta-Ethics").

committed to them. The strategy, then, will be to argue that there are certain principles or norms that are *constitutive* of having certain features within one's practical standpoint at all. Here, the mark of a constitutivist claim is that, if one is already committed to X, then one *just is* committed to Y as well – such that if one attempted to deny a commitment to Y, one would in that same move be denying a commitment to X (which one cannot do, insofar as it is a secure premise that one is already committed to X).[22] For instance, one could therefore attempt an argument where, if someone is *already* committed to valuing *anything at all* ('X'), then, in virtue of that, she will also be constitutively committed to some stronger moral principle ('Y').[23] So, according to our picture, there is ultimately normative objectivity *by our own lights*: in bestowing normativity upon the world through our practical standpoints, we cannot help but also bestow certain deeper principles or standards that have a binding force on us.

Now, constitutivist explanation *is* a form of explanation, and it therefore does not make an appeal to the bruteness of normative facts. When asked what justifies some normative principle, and what makes it binding upon the individual, the theorist can respond with an argument to the effect that the individual, in virtue of something that she is *already* doing in her practical standpoint, has normatively committed herself to this principle. Of course, the success of such a view will rest on whether the particular arguments are valid and sound, and so the proof must lie in the pudding – but at least the *form* of the strategy promises to avoid brute normative facts. In Chapter 3, I shall argue that it is constitutive of normative deliberation in the present that we commit ourselves to the normative authority of an objective range of counterfactual endorsements that are possible. And in Chapter 5, I argue that the human practical standpoint contains a commitment to the principle that valuing *qua* valuing is a sufficient condition for love – which means we are committed to the

[22] See also Street (2008: 228).
[23] We will attempt an argument along these lines (§5.2).

idea that *all* normatively conscious perspectives matter, i.e. that all perspectives are deserving of love.

Discussion of objectivity often comes along with the notion of *universality* as well. We shall be interested in both notions, but we must also keep them distinct. To speak of normative *objectivity*, in our sense, is to speak about how some agent can be beholden to certain normative standards that transcend her particular thoughts or other attitudes, and of which she can be seriously mistaken. Of course, it is also possible for someone to use the term 'objective' in a stronger sense, in which it is simply equated with the idea that there are facts or properties or entities 'out there' – we will not be interested in that stronger sense of objectivity, which is simply synonymous with realism.[24] *Universality*, on the other hand, is the idea that the *same* normative standards apply to all sorts of different individuals. This is why the two concepts are not the same: just because some standard is *objective* for Kevin, in that it continues to apply to him, independent of his changing thoughts and attitudes, does not mean that it is also *universal* in the sense that it is also the *same* standard that binds *Albert*. With that said, we can state that we shall be interested in vindicating both universality as well as objectivity, for with universality, we will be making arguments about a practical standpoint that characterizes the human being in general, across different cultures and time.[25] Whether every conceivable alien will also fall under the umbrella of these arguments will depend on whether some alien shares the features of the practical standpoint that we take to be constitutive of these deeper principles, such as a creature's valuing with certain capacities (see §5.2).[26]

---

[24] See Adrian Moore for a distinction between *world-directed* objectivity and *conclusion-directed* objectivity (2003: 5), as well as Street for a distinction between *realist objectivity* and *constructivist objectivity* (2016a: 173), with our interest being in the latter of these distinctions.

[25] Or, at least the human being in her focal case. For instance, someone in a permanently vegetative state is still a human being, but not necessarily capable of the deliberation and action that would make her an *agent*, and thereby covered by the scope of the arguments that I offer. This is what I mean by focusing on the focal case of being a human – human, but with the relevant capacities that I shall discuss as the theory develops.

[26] This might allow us to navigate a middle position between so-called Humean and Kantian constructivists (see Street 2012). Like the Humean constructivist, we want to describe certain practical standpoints that could have been otherwise, in that we can conceive of practical

Now, I want to make one further point of clarification: while I have emphasized that we want to develop a view that avoids all *brute* claims about normativity, this does not mean that we think that there is no philosophical bedrock *anywhere*. For instance, if we explain normativity as arising from the mind, someone could still ask about what explains the existence of minds – and that would be a legitimate question, though not one that I want to take a position on in this dissertation. My point is simply that, just because we are setting aside the idea of brute *normative* truths, does not mean that we think there are no fundamental layers of explanation. It is the distinctiveness of claiming that *normativity* involves brute facts that incurs the sort of objections that we reviewed earlier, and those discontents with *normative* realism are what motivate our project – but philosophical explanation still has to bottom out somewhere, and for our project, we will simply treat the mind as something that we can take for granted, as something that does exist and which we are familiar with, though we leave open the further question of how to explain the existence of normative consciousness.[27]

What is the promise of such a theory, then? It is a theory that promises to vindicate normative objectivity, while avoiding incurring the same discontents that realism faces. Furthermore, return again to the illustration of normative disagreement that we imagined – importantly, our theory will offer a different diagnosis about how such disagreement comes to an end. Suppose again that there is a pure normative disagreement between you and someone else about what is good, beautiful, or right. Rather than insist that the other person has faulty access to some mind-independent normative truth, you can attempt to *reason with* that person from her starting perspective, and to offer her a series of

---

standpoints that are less sophisticated than the ones we will be focusing on; but like the Kantian constructivist, we will want this to be quite a wide net, which captures the focal case of being a human being. There is also a delicate relation here to some of the points made in Street (2009a) – while I agree with her that there are certain conceivable creatures that are not subject to moral requirements, the class of creatures that I shall target is wider than what she outlines in that article, and it includes Caligula (see §5.2).

[27] I also do not want to take a stance on whether the phenomenon of the mind is reducible to physical facts or not, or how the mind might supervene on them; it is enough for our purposes that phenomenal consciousness exists, and that we know what it is like to occupy a practical standpoint on the world.

arguments in dialectic, so as to show how her existing commitments bind her to certain moral and normative standards of which she is not yet aware.[28] And at every point in this dialectic, there will be no mere *insistence* about the correctness of some brute normative truth, but always an explanation for why that normative truth applies to one's interlocutor, and where its existence comes from. This is not to say that we can convince everyone, in the sense that they will always listen to or agree with us – but it at least promises an answer in dialectic that does not end simply in a display of confidence.

1.4 Notes on Existing Literature

Having outlined our project, I want to make some further comments on its relation to some existing literature. I have mentioned that our project can be called an *antirealist* position, in that it avoids the presuppositions of realism; and it can also be called a *constructivist* theory, in that it maintains that "normative truth consists in what is entailed from within the practical point of view" (Street 2010: 367).[29] However, I want to emphasize that both of these labels come with certain unhelpful connotations. *Antirealism*, for instance, connotes a primarily negative project, whereas we are hoping to provide an account that *vindicates* the objectivity of normativity; and on that score, antirealism also suffers from a heavy association with relativism.[30] The term *constructivism*, though it fares better on these fronts, also suffers from too much of a metaphor about *construction* on the part of the agent. As my reader shall come to see, I will be arguing that much of our experience of normativity is *involuntary*, including our inescapable commitment to certain objective, binding normative standards that

[28] This was a point that Williams suggested when he spoke of reasoning from another's perspective (1995: 36), and this requirement about reasons is also expressed thoughtfully in the recent work of Fan (2021).

[29] Our project is what Street furthermore calls a *thoroughgoing* or *metaethical* form of constructivism, which does not bring in any substantive ends into its starting characterization of the practical standpoint (2010: 367-369).

[30] 'Antirealism' is further made complicated by the fact that some theorists who reject mind-independent normative facts nevertheless refer to themselves as *realists* in some thinner sense – see Lewis (1989) and Korsgaard (1996) as examples.

arise from our practical standpoints – and so there is some sense in which we do not actively *construct* these norms for ourselves, but rather come to find ourselves as bound by them. Even Rawls, who is responsible for introducing the term constructivism into the recent literature, wrote that "not everything, then, is constructed; we must have some material, as it were, from which to begin" (1993: 104). I shall therefore prefer to refer to our project as developing a *mind-given* theory of normativity.

I also want to emphasize that there are numerous views in the existing literature which might sound similar to our project, but which are not in fact aligned with our presuppositions. For instance, there are many views that call themselves 'subjectivist' or 'internalist,' or which advocate 'procedural' or 'idealizing' or 'non-alienating' requirements rather than 'substantive' ones, but which do *not* ultimately avoid the existence of brute and mind-independent normative facts. This point is subtle, and often lost track of in discussion, but it is nevertheless crucial. As an example, consider that there are some influential views which maintain that there must be a *necessary* condition between our normative attitudes and certain normative truths, but which do not make the stronger claim that normative truths are *explained by* or *come from* our normative attitudes. Sample, for instance, the following statements from the literature:

> What is intrinsically valuable for a person must have a *connection* with what he would find in some degree compelling or attractive, at least if he were rational and aware. (Railton 1986: 6, emphasis added)

> An object, event, state of affairs, etc., ϕ is good for an agent x *only if* x takes a valuing attitude (of the right sort) towards ϕ. (Dorsey 2017a: 687, emphasis edited)

> Subjectivists about welfare claim that something is basically good for you *if and only if (and to the extent that)*, in a certain familiar sense, it *suits* or *fits* you. (Lin 2017: 354, emphasis added)

> Existence internalism is the idea that there is a *necessary connection* between something's having a certain positive normative status (its being right or its being intrinsically prudentially good) and its ability to elicit positive psychological responses from human subjects. (Hawkins 2019: 95, emphasis added)

These statements maintain that there must be at least a necessary connection between our normative attitudes and certain normative truths, such as what is good for us.[31] But notice, now, that all of these statements could be understood as compatible with the existence of brute and mind-independent normative facts: for instance, these stipulations could *themselves* be considered brute facts about normativity, if we are not given some deeper explanation of what *makes* such claims true. Dale Dorsey insightfully makes this point: a 'resonance constraint' of the sort described above is "compatible with the independent existence of value facts" (2017a: 694).[32] This can also be understood as the point that necessary *covariation* between our normative attitudes and normative truths is not the same as our normative attitudes being what *explain* or *ground* those normative truths, in the sense of being the ultimate source for them. As Jonathan Dancy notes, "the 'grounding' or 'providing' relation is different from, stronger than, and much more interesting than the much weaker relation of 'being a necessary condition for the existence of'" (2000: 17) – and Dancy notes that this is a point that sometimes escapes Parfit's discussion.[33] Or, as Selim Berker puts it, "grounding is not counterfactual covariation" (2020: 197). On our theory, we are making a stronger claim about what the phenomenon of normativity *is*, and where it *comes from*, rather than just offering certain necessary constraints on

[31] Cf. Sobel for an account of subjectivism that makes the grounding claim, rather than the covariation claim: "subjectivists claim that the relevant sort of desire grounds, not merely tracks, the truth of claims about what is good for a person" (Sobel 2009: 337).
[32] See also Lin (2017: 355).
[33] Dancy is speaking here of Parfit (1997), though I believe the issue still persists into Parfit's *On What Matters* (2011). For instance, it may be seen in Parfit's introduction of subjective theories in Part One (2011a: 58-61), though cf. Parfit (2011b: 289-290).

normativity – we are claiming that normativity is a phenomenon that exists within, and arises from, the practical standpoint.

Continuing onwards, I also want to emphasize that some views which advocate 'procedural' or 'idealizing' requirements upon reasoning are still ones that can involve a commitment to the existence of brute and mind-independent normative facts. Julia Markovits (2014) is a primary example of this, as we shall discuss at greater length in Chapter 9.[34] Markovits maintains that there are certain procedural requirements upon reasoning that contrast with *substantive* requirements, i.e. ones which directly dictate what ends we should hold (2014: 51-52). But if we ask Markovits what justifies the *normativity* of her procedural requirements, she admits at subtle points that these requirements are "irreducibly normative" (2014: 55) and even emphasizes that "we are in no way the authors of the requirements of procedural rationality, or the sources of their normativity" (2014: 73n4). In this way, Markovits is committed to the existence of brute and mind-independent normative facts – and anytime some theorist advocates a procedural or idealizing requirement upon reasoning, we can similarly ask them where the normativity of *that* requirement comes from. Insofar as the theorist claims that its normativity is irreducible or brute, then that theorist will be committed to realism, just like the more familiar realists we have mentioned – importantly, note that even *one* claim about a brute normative fact is what moves a theory into the territory of realism, for it is any such assertion about bruteness that incurs the sort of dogmatic, metaphysical, epistemological, and other objections that we have reviewed.

The general lesson is that, within the recent literature, it can behoove us to be careful in understanding what a view is really claiming, even if it comes under the heading of some label that might sound related to our project.[35] Just because a view *calls* itself internalist or otherwise does not mean that it has the

[34] See also Chapter 8 on Smith.

[35] For similar reasons, I would also distance our project from views that maintain that one individual's reasons straightforwardly depend on the normative judgments or attitudes of *another* person (see Schafer 2014 and Manne 2016), for such a view could leave open or unanswered the question of why I should care about another's judgments about my reasons.

same departure point as we do, as setting aside *all* brute and mind-independent normative facts. I believe that, among the recent literature, it is Korsgaard and Street who have most clearly attempted projects along the same starting presuppositions as ours, and in that sense they are most closely aligned with the parameters of our inquiry.

1.5 The Project Ahead

The chapters ahead develop this theory of normativity. We start with the understanding that normativity is rooted in the practical standpoint, the standpoint of conscious *being* – and we move from there to build a more robust account of normativity. These chapters not only attempt to vindicate objectivity in both the moral and non-moral domains, but they also address how our faculties render the world into a normative place, confront certain questions about nihilism, and engage with other prominent positions from the recent literature.

The first half of the dissertation, running from Chapters 2 through 5, develops the positive vision for this theory of normativity. Chapter 2 begins with a preliminary and fundamental topic: that of how the human mind renders the world into a normative place. I argue against a widespread thesis in the recent literature that I call Attitude Monism: the idea that we can identify one *single* attitude that is an exhaustive source of normativity. In lieu of that, I articulate a multi-faceted account of the sources of normativity, culminating in the notion of a normative *endorsement*, which refers to our most internally authoritative normative stance on some matter. I also introduce the notion of a normative *lens*, which describes the background conditions which shape the particular normative attitudes that we experience.

Chapter 3 then builds a constitutivist account of how there is normative objectivity in the pre-moral domain. I argue that we are all constitutively committed to the authority of our strongest counterfactual endorsements, which outsources normativity to an objective hierarchy of mind-world potentialities;

and I also argue that our normative endorsements create further standards by which we can make normative mistakes. These points promise an important structure of normative objectivity, where there are normative standards that obtain independently of our present thoughts or attitudes.

Chapter 4 addresses a series of nihilistic concerns for a mind-given picture of normativity. Confronting the concern that all value and beauty will fade alongside the disappearance of sentient minds from the world, I contend that reflection on potentialities can allow us to recover a set of unchanging and everlasting facts that mirrors the realist's mind-independent values. I also argue, against certain objections, that mind-given normativity is both *genuine* and *real.*

Chapter 5 culminates our positive project, in its attempt to develop an original account of *moral* objectivity. I focus on the idea that there is a sufficient condition that is embedded in the way we value: that valuing *qua* valuing that is a sufficient condition for the deservingness of love. Offering two independent arguments for this idea, I contend that we are all under the moral requirement to maintain that the valuing of every perspective *matters* in the same way we take our own concerns to matter – and I call this conclusion the **Principle of Universal Love**. We must therefore take the valuing of other perspectives to count as binding sources of reasons for ourselves; in other words, we are all inescapably committed to the conclusion that every point of view is deserving of love.

The second half of the dissertation, running from Chapters 6 through 10, turns to examine prominent accounts from the existing literature that seek to vindicate objective moral requirements, on some kind of constitutivist, constructivist, or internalist basis. Chapter 6 focuses on Korsgaard's *The Sources of Normativity* (1996), and Chapter 7 discusses her later work in *Self-Constitution* (2009). Chapter 8 then examines Michael Smith's recent constitutivist position (2013 to 2021), and Chapter 9 examines Julia Markovits' Kantian, internalist position – and I suggest that Smith and Markovits do not thoroughly reject the existence of brute and mind-independent normative facts. Finally, Chapter 10 engages with Street's recent Buddhist-inspired

constructivism (2016). I examine these views in both a critical and constructive fashion, arguing that while they are each inventive and important in their own right, they each also face some important questions or objections.

What emerges from these pages is the development of an original theory of normativity (Chapters 1-5), along with a novel examination of some recent arguments in the literature (Chapters 6-10).

2 The Normative World

2.1 Introduction

We are beginning with the proposition that normativity arises from the standpoint of conscious *being* – from what I have described as the practical standpoint (§1.3). Our aim is now to develop this theory of normativity, in order to see whether we can render it holistically plausible and attractive. This means that we must address a preliminary and fundamental question: from the starting presuppositions of a mind-given picture, how exactly does the human mind come to render the world into a normative place?

This is a question that is both inherently important for a mind-given theory of normativity, and it also requires navigating through some disagreement in the literature. Much of the literature is guided by an implicit assumption that I call *Attitude Monism*: the thesis that we can identify *one* faculty or attitude that serves as the exhaustive grounds of normativity.[36] One of my initial aims in this Chapter will be to argue that Attitude Monism is mistaken, before then turning to develop a positive account of how our practical standpoints constitute normative truths.

I start by identifying and arguing against the widespread assumption I have called Attitude Monism (§2.2). I then develop my own account of how our attitudes constitute normative truths, which puts at the center a notion that I call normative *endorsement*, which refers not to a specific attitude we can exercise, but to our more general capacity to take an internally authoritative stance on some normative matter *through* different attitudes (§2.3).[37] Following that, I discuss some of the theoretical upshots of the picture I have advanced (§2.4),

[36] I use 'faculties' and 'attitudes' as interchangeable. Also, I use the term 'grounds' here in a loose sense, to encompass theses that concern a co-variating connection, as well as a stronger source claim (cf. §1.4).

[37] Will my emphasis on normative endorsements prove to be a form of Attitude Monism? No, because a normative endorsement does not refer to a *specific* psychological attitude, but is rather used as a philosophical term of art, to capture the idea of internal authoritativeness *through* different possible attitudes.

and I close by formulating the notion of a *normative lens*, which concerns the background conditions that color our experience of normativity (§2.5).[38]

2.2 Against Attitude Monism

When we ask about how the practical standpoint gives rise to normativity, a natural place to begin is with our faculties or attitudes. There is a widespread assumption in the existing literature that we must select *one* attitude that serves as the exhaustive grounds of some relevant set of normative truths, e.g. our normative reasons or prudential goodness. I call this assumption *Attitude Monism*. In this section, I want to identify what an influential assumption Attitude Monism is, and I shall then argue against it, on account of both substantive and theoretical considerations.

Attitude Monism is a position that can attach to any type of attitude, whether that is desire, second-order desire, belief or judgment, choice, or pleasure. The claim of the Attitude Monist is that, once we identify which one of these attitudes is the correct one, we can then use it as an exhaustive ground of our normative reasons or prudential goodness. The most well-known version of Attitude Monism attaches to the attitude of *desire*,[39] in the form of influential 'Humean,' 'internalist,' or 'subjectivist' theses:

> Since Internalist theories are desire-based… all reasons to have some desire must derive from other desires. (Parfit 1997: 128)

> Humeans hold that all normative reasons have their source or ground in the agent's desires. (Darwall 2003: 437)

> On a widely accepted story of human agency… all practical reasons are grounded in the present desires of the agent. (Chang 2004: 56)

[38] None of my remarks in this Chapter should be construed as seeking to vindicate normative objectivity yet – a task that will occupy Chapters 3 and 5.
[39] Note that I use the term 'desire' on its own to refer to *first-order* desire, and I also exclude usages in which it refers to a broader set of attitudes (e.g. Parfit 1984: 117 and Smith 1987: 54).

> Most philosophical discussions of the Humean Theory of Reasons… assume that to be a Humean, you have to take a view about what kind of psychological feature explains the difference between Ronnie and Bradley. They assume that a Humean must hold that this psychological state must be a *desire*. (Schroeder 2007: 6)

> What defines an account of well-being as subjective is the thesis that an agent's desires ground what makes something good for her. (Sobel 2009: 336)

Attitude Monism about desire is so influential that it can sometimes be seen as a *default* position for antirealist or subjectivist theories: as Street notes, "it is widely assumed that if an agent's reasons are always a function of his or her attitudes, then the attitudes in question must be the agent's *desires*" (2012: 42).

But Attitude Monism can also attach to other attitudes, and such views have been increasing in light of the recognition that desires are implausible as an exclusive source of normative truths.[40] For instance, Peter Railton maintains that prudential goodness consists in the attainment of second-order desires: "an individual's *intrinsic* good consists in attainment of what he would in idealized circumstances *want to want* for its own sake" (2003: 54, emphasis added). David Lewis maintains a similar view about second-order desires as what explain values in general, claiming that "valuing is just desiring to desire" (1989: 163). Dale Dorsey, by contrast, has argued against Attitude Monism about desires and second-order desires, and he has instead advanced a version of Attitude Monism about beliefs and judgments: "instead of suggesting that x values ϕ to the extent that x desires ϕ, subjectivists should say that x values ϕ to the extent that x *judges* that ϕ is good for x" (2012: 408, emphasis added).[41] And Korsgaard can be read

[40] For instance, Harry Frankfurt writes that desires constitute "a more or less standard but nonetheless rather meager conceptual repertoire," noting that they are "heavily overburdened, and a bit limp" (2004: 10). See also Street (2012: 42-44), Sobel (2016: 3-4), and Scanlon (2021: 32).

[41] Note that Lewis and Dorsey are using the terms 'valuing' and 'values' above in a general sense, almost like a placeholder notion for the more specific attitude to be identified in their theories.

as an Attitude Monist about choice (1996: 163), as I discuss later (§2.4). To appreciate how ingrained the assumption of Attitude Monism is in the literature, consider how Eden Lin recently presents the "three obvious candidates" for a subjectivist account of prudential goodness:

> Judgment Subjectivism: x is basically good for you at W if and *only if* under the right conditions, you *believe* that x is basically good for you at W.
>
> Desire Subjectivism: x is basically good for you at W if and *only if* under the right conditions, you *desire* X.
>
> Value Subjectivism: x is basically good for you at W if and *only if* under the right conditions, you *value* X.[42]

Each of these options are Attitude Monist positions, because each one stipulates *one* attitude as a necessary condition for prudential goodness – implying that it is *either* judgment that constitutes prudential goodness, *or* desire that constitutes prudential goodness, *or* valuing that constitutes prudential goodness, but not some combination of these attitudes.[43]

Let me now argue that Attitude Monism is mistaken, on account of both substantive and theoretical considerations. First, at a substantive level, any kind of Attitude Monism will seem to leave out some other plausible grounds of our normative reasons or prudential goodness. For instance, consider a case against Attitude Monism about *desire*:

> **Burnout.** Because Albert is burnt out from being overworked, he has very few first-order desires at all – for instance, he has no desire to take care of his own health, or to do the things that would make him feel more happy, like seeing

[42] See Lin (2017: 356, emphasis added).
[43] Valerie Tiberius might be read as an Attitude Monist about *valuing*, an attitude that involves aspects of judgment, reflective endorsement, and emotion (2018: 13). Although Street can be read similarly at particular points in her writing (see 2012: 42), she otherwise seems committed to a form of Attitude Pluralism (see 2016a: 176).

> friends or reading for leisure. Albert does *believe* that it would be good if he did these things, and he even has a *second-order desire* to have the first-order desires to do these things; he simply lacks the first-order desires.[44]

It seems implausible to claim that Albert has no normative reason to take care of his own health, or that doing so is not part of his good, simply because he lacks the first-order desires to do so. Not only does taking care of his health seem like a part of his good, but moreover, he even has the beliefs and second-order desires that would support those conclusions.

Next, consider a case that challenges Attitude Monism about *second-order desire*:

> **Parched.** Jane is parched, because she has not been drinking water, and she has a strong first-order desire to drink some water. However, she is on the verge of having a heat stroke, with diminished experiential and mental faculties, and she fails to have a second-order desire to desire to drink water.

It seems implausible to claim that Jane has no normative reason to drink water, or that drinking water is not part of her good, simply because she lacks a second-order desire to do so. Indeed, from the perspective of a mind-given theory, it seems like her first-order desire to drink water is sufficient to constitute some presumptive reason for her to drink water.

Continuing onwards, consider a case that challenges Attitude Monism about *belief* or *judgment*:

> **Abuse.** Someone who has been the victim of significant physical and emotional abuse feels a strong desire to no longer be abused. However, because the abuse has led to confusing patterns of thought and affected her capacity for independent judgment, she fails to form a belief or judgment about what would be good for her.

[44] See Dorsey (2012: 410) for a different case of burnout. This point can also be made through cases of depression – see Tully (2017) and Spaid (2021) for discussion.

It seems implausible to claim that it is not part of this person's good that the abuse comes to an end, simply because she is not at a place where she can form a judgment or belief about the matter.[45]

Finally, consider a case that challenges Attitude Monism about *choice*:

> **Sunshine.** You have stepped into the sunshine and breeze, and you immediately come to enjoy the feeling of this atmosphere on your senses. However, in the first moments of this feeling, you have not yet had the time or thought to reflectively approve of the sunshine as good from the standpoint of choice; you are simply focused on enjoying it.

It seems implausible to claim that, in cases where we are having simple enjoyments through our senses, these experiences are not yet good until we have further stamped them with the approval of reflective choice. From the perspective of a mind-given theory, it seems that the enjoyment, even prior to its being chosen, is what makes the experience something that is good from the person's perspective.

Cases like these suggest that Attitude Monism is an implausible position at the substantive level, regardless of which specific attitude it attaches to. This is because it seems like there are always cases in which we have some reason, or there is some thing that is part of our good, but where these reasons or goods cannot be exhaustively explained through appeal to one single attitude.

But for the purposes of our theory, there is also a second reason to reject Attitude Monism, which is theoretical: Attitude Monism does not follow from the presuppositions of a mind-given view. According to a mind-given (or antirealist) theory of normativity, normativity arises from within the practical standpoint. But given this metaphysical picture of the general source of normativity, there is no reason to think that normativity must be limited to the exercise of one attitude: for it is possible to engage with the world *normatively*

[45] See also Lin's argument that newborn babies cannot have beliefs or judgments, yet still have a good (2017).

through many different attitudes. For instance, we can *enjoy* certain feelings, *desire* certain foods, *value* another's love, *judge* certain pieces of art to be beautiful, or *choose* to commit to certain projects; and we can *dislike* certain sensations, *judge* that certain artworks are bad, *desire* to be in a greater state of freedom, and experience a wide array of normatively valenced emotions. Each of these attitudes involves some experience of the world that has a normative dimension to it from within the practical standpoint, and it would hence be *ad hoc* to limit normativity to just one attitude at the outset. For instance, from the perspective of a mind-given theory, judgments can be one way in which we take the world to be normative, even if we do not have a corresponding desire:

> **Exercise**. I take exercise to be good for me, through a judgment or belief. However, I do not experience much desire to exercise.

Here, the judgment *involves* representing something as good from within the practical standpoint – and therefore bestowing some goodness on it. Or, for instance, a mind-given theory can also say that pleasure is a source of normativity through which we find things to be good, even if we do not have a corresponding belief or choice:

> **Taste.** The taste of some new vegetarian dish leaves a delicious flavor on my tongue that I greatly enjoy. However, I have not had time to form a belief or reflective choice about the matter yet.

Because the exercise of our various attitudes can each involve some normative relation to the world from within the practical standpoint, Attitude Monism about any particular attitude will prove to be *ad hoc* from the basis of a mind-given theory. Since we are developing a mind-given theory, the fact that a number of our attitudes involve *normative* experiences is a theoretical reason to be against Attitude Monism.

For these substantive and theoretical reasons, I believe we should reject Attitude Monism.

2.3 The Positive Picture

If we have set aside Attitude Monism, then we should accept some form of Attitude Pluralism: the thesis that a range of attitudes can constitute normative truths. Attitude Pluralism is not a new position in the literature, and we can read Williams and Street as having maintained some form of Attitude Pluralism.[46] However, emphasis on Attitude Pluralism is not often made explicit in the literature, given the intramural debates I have surveyed between Attitude Monists, and moreover, the acceptance of Attitude Pluralism still leaves numerous questions open, including what the relation *between* our different attitudes are, and how any of them could be ultimately authoritative.

My aim now is to articulate a positive account of how the practical standpoint gives rise to normative truths. I shall maintain that a range of our basic attitudes, like enjoyments and aversions and desires, provide an initial layer of normativity to the world. But we humans then have the capacity to set more authoritative normative directions for ourselves, through more sophisticated faculties like judgment, choice, and valuing. After discussing these faculties, I then formulate a concept that I call normative *endorsement*, which refers to the most internally authoritative stance that one can take on a normative matter.

Let me begin with the point that our basic enjoyments, aversions, and desires provide the world with an initial layer of normativity. This is due to a simple reason: our experience of these attitudes involves our representing the world *as good* or *as bad* to ourselves, from within the practical standpoint. According to a mind-given theory, that simply is what it is for there to be

[46] Bernard Williams held that internal reasons could be grounded in "such things as dispositions of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called, embodying commitments of the agent" (1981: 105), noting that this set of attitudes "should be more liberal than some theorists have been about the possible elements" in it (1981: 105). Street uses 'valuing' as "a philosophical term of art meant to encompass a wide range of attitudes and experiential states" (2016a: 176). Some theorists who discuss 'pro-attitudes' can also be understood as Attitude Pluralists.

normativity: for us to relate to the world in a normative dimension from within the practical standpoint. Consider what it is like to *enjoy* the feeling of sunshine lapping on one's skin: to *like* that feeling is to experience that sensation as irreducibly good. Similarly, to *dislike* the feeling of pain is to experience that sensation as irreducibly bad. Not only do we use these normative concepts of goodness and badness in an irreducible manner to describe these experiences, but it is also *those experiences themselves* that have an irreducibly normative dimension to them, in the sense that they cannot be reduced to mere neutral description: there is something fundamentally valenced about these experiences, in a positive or negative sense, from within the practical standpoint.[47] These claims do not mean that this experience must point to certain standards of normativity that are external to the mind: they just mean that, from within the practical standpoint, the sensations are experienced in irreducibly normative ways, and then expressed through irreducibly normative concepts. It is these kinds of basic enjoyments and aversions that we share with some of the other animals: although they may not have the same conceptual schemes we have, their basic enjoyments and aversions also constitute, from their practical standpoints, irreducibly normative encounters with the world.[48]

Although desires are somewhat more controversial, I believe that they can provide an initial layer of normativity in the same way. This is because, when we have desires, we usually represent the imagined object of desire with some kind of positive valence.[49] Return again to the example of Jane, who desires to drink some water: when she has that desire, she represents the drinking of water *as good* from within her practical standpoint, through that positive valence, and she therefore provides some normativity to the drinking of water. The same point applies to a wide range of examples, such as when we desire to

[47] See §4.3 for further elucidation of the idea that our normative concepts and experiences have an irreducibly normative dimension to them, and why this is different than the realist's claim that there are irreducibly normative *facts*.
[48] Korsgaard insightfully formulates some of these points in *Fellow Creatures* (2018).
[49] For instance, Scanlon writes that "desiring something involves having a tendency to see something good or desirable about it" (1998: 38), and Dancy writes that "most of the things we want, we want for a reason. This reason will usually be some good that we see in the thing that we want" (2000: 35).

eat certain foods, desire to listen to particular pieces of music, desire to improve ourselves, or even desire to be in communion with the divine: in desiring these things, we represent the object of desire *as good* through our ascription of some positive valence upon it, however simple or complex. There are some delicate questions here about whether *all* desires necessarily involve the representation of such a positive valence, and whether such a representation is intrinsic to the desire itself or is something that comes along with the desire – questions which I discuss at greater length elsewhere.[50] We can proceed simply with the claim that desires *usually* come along with the representation of their imagined objects with a positive valence, and that constitutes some goodness from within the practical standpoint.[51]

We can claim that these basic encounters with goodness and badness hold some constitutive connection to there being a *pro tanto* normative reason in relation to that goodness or badness. By this, I mean that when we ground some goodness or badness from within the practical standpoint, we likewise ground a reason of some kind that constitutively or conceptually follows from there being that goodness or badness. Although this might seem like a controversial point, it need not be upon examination.[52] A normative reason means something like 'a consideration that counts in favor of' – and when we

[50] See my *Prolegomena to an Internalist Meta-Ethics* (2021: §2.4). One conceptual choice here is whether we want to distinguish certain urges, whims, or impulses – which do not involve the representation of their objects as having a positive valence – as fundamentally different from desires. For instance, think of examples in which one wants to drink a can of paint (Davidson 1963: 686), or to turn the knobs of radios without actually hearing the music or news (Quinn 1993: 236): insofar as we treat these as mere urges and not desires *qua* desires, we might preserve the conceptual claim that *all* desires involve the representation of a positive valence. Dancy (2000: 36) and Shemmer (2007: 328-329) draw this distinction between urges and desires.

[51] Furthermore, note that what we ascribe a positive valence upon is the *imagined* object of desire, which can diverge from real objects in the world – a point that we shall return to later when we discuss objectivity (see §3.3). This is why a desire for a gin and tonic does not ground a reason to drink a clear liquid in front of me that is *in fact* petrol. See also Sobel (2009: 345) and Dorsey (2017b: 207-208).

[52] The idea that normative reasons are grounded in conceptually prior claims about value or goodness can be found, in different forms, in Raz (1999: 22-23), Crisp (2006: 37 & 61), and Maguire (2016: 237). As Raz writes, "reason is then explained in part by invoking value" (1999: 22). See Kiesewetter (2022) for criticism of this idea. Some other theorists – who champion a 'buck-passing' account of goodness (Scanlon 1998) – maintain that the notion of a reason is the conceptually basic one. My sympathies lie against buck-passing views, but I will not defend that point here; see Crisp (2005) for discussion.

ground something *as good* or *as bad* from within the practical standpoint, we are constitutively implying that this goodness or badness *counts in favor* of *something*. Perhaps this claim can seem more significant than it really is, because some authors have operated with a rather restricted notion of a reason, such as only a reason to *promote*; but there can in fact be a wide range of reasons. For example, when we ground something as good, that can conceptually imply reasons to admire, to stand still and appreciate, to cherish, to love, to acquire, to respect, to honor, to recognize, and so on; and when we ground something as bad, that can conceptually imply reasons to avoid, to criticize, to remove, to reduce, to withdraw from, and so on.[53] Moreover, these are *pro tanto* reasons – there being a 'consideration that counts in favor of' *to at least some extent* – but which may not turn out to be decisive. The point is that there would be a constitutive or conceptual mistake in claiming that the feeling of sunshine is good, but that there is no reason at all to stand in it; as there would be, if someone claimed that pain is bad, but that there is no reason at all to avoid it.[54]

I have argued that our basic enjoyments, aversions, and desires can provide an initial layer of normativity to the world, and therefore pro tanto reasons. But humans are also capable of judgments or beliefs, choices, and valuing. I now want to claim that it is through our exercise of these more sophisticated attitudes that we can set *more authoritative* normative directions for ourselves, which do not *extinguish* the normativity of our enjoyments, aversions, and desires, but which can *overrule* or *strengthen* their normative force.

Consider first what it is like to render a normative judgment. Suppose that the pleasure of relaxation is tempting one to procrastinate on one's work, but one judges that the value of accomplishing some aim – say, of completing a

[53] For the point that there are a diversity of reasons, see Rabinowicz & Rønnow-Rasmussen (2000: 46) and Frick (2020: 64), among others. Cf. Maguire (2018) for a recent argument that there are no reasons for affective attitudes.

[54] This might instead be put as the claim that goodness constitutively entails certain *conditional* reasons, such that *if* it were possible that we could interact with the object that is good, *then* we would have some reason in light of it. This might make sense of examples where something might be good or beautiful, but too far out of our reach – see Moore's example of a universe being good on its own (1912: 28), as well as my discussion about value in §4.2.

philosophical project – matters more than the temporary enjoyment of that relaxation. What is happening here from within the practical standpoint? Here, we acknowledge that there is some *pro tanto* good or reason that is arising from the pleasure of relaxation, but we set a *more internally authoritative* direction for ourselves through our judgment that the philosophical project matters more. Through our judgment, we exercise our capacity to weigh the relative importance of different normative goods to us, and to assess one as more significant than the other. The reason that our reflective judgment constitutes a stronger reason to pursue the philosophical project, and constitutes a greater goodness to it from our perspective, is that this simply is *what it is* to have a stronger reason, on our mind-given theory of normativity: it is to take something to be more reason-giving from within the practical standpoint. We can see such an authoritative role for judgment in a range of examples, such as when we judge that it is worthwhile to suffer for some cause, to sacrifice our time for someone else, to enlist in a war, or to pursue some challenging vocation in the service of some aim. Or indeed, when we judge that some desire is *good* – for instance, the desire to improve ourselves – we can add further fuel to the normative import of that desire, thereby strengthening the reasons that it provides.

Choice is perhaps the even more paradigmatic faculty through which we can, from within the practical standpoint, set an authoritative direction for ourselves, in spite of desires, pleasures, or pains that beckon towards the opposite direction. When we make a choice from within the practical standpoint, we represent ourselves as exercising a *will* to put ourselves behind some particular path, for some chosen reasons.[55] And not only can such choice be used to overrule desires or pains from within the practical standpoint, but they can even be used to overrule *judgments*. For instance, imagine a case in which you *judge* that some item in the shop is beautiful or valuable, but you nevertheless

[55] Freedom of the will, and the very notion of exercising our will, should be understood in a practical, phenomenological sense here: this is *what it is like* to exercise our will from within the practical standpoint, on the *assumption* that we are free. I want to take no further stance on metaphysical questions surrounding a free will – and in this way, this feature of the view is to some extent Kantian.

choose not to purchase it – that would be an example of our exercising a more authoritative attitude from within the practical standpoint. For this reason, choice can be even more authoritative than judgment.[56] But it is important to emphasize that, when our choices set stronger reasons for ourselves, that are more authoritative than the ones we already had, the *mind-given* claim is not that we create some *further* metaphysical entity (a reason) that is then distinct from the will. Rather, on our view, when we set a more internally authoritative direction for ourselves, we *just are* representing a stronger normative reason from within the practical standpoint, through our choice – without creating anything that exists beyond the practical standpoint.[57]

There is also the notion of *valuing* in the literature, which, when used to refer to a specific attitude of its own, can be distinguished with its emphasis on certain patterns of emotional experience. For instance, Tiberius describes valuing as involving "a relatively stable pattern of emotional, motivational, and cognitive dispositions or tendencies towards what is valued" (2018: 11).[58] In a similar vein, Street has described an attitude of valuing as distinct from desire, whereby valuing "is characterized by all the range, nuance, and depth of human emotion and feeling," including by its "greater structural complexity" (2012: 44).[59] Such valuing can also override particular desires or enjoyments. For instance, we might *desire* to run away from certain hard difficult interpersonal

[56] If we make a *holistic* judgment about a situation, that might end up looking synonymous with our making a *choice* about a situation. Note that I am using the concept of 'choice' here in a way that is divorceable from what we actually end up doing, such that one can choose to pursue path X, but still fail to follow through on it.

[57] Chang contrasts 'constitutive ground' with metaphysical 'creation': "An act of will can be the constitutive ground of something's being a reason and *it can also be what creates, as opposed to constitutes*, the fact that it is a reason" (2013: 168, emphasis added). We are interested in something closer to the former: nothing new is metaphysically created, but rather, our making a choice *just is* our taking ourselves to have an authoritative reason *from within the practical standpoint.* Nevertheless, as I shall later emphasize, this does not make the normative reason, or normativity, any less *real* (§4.4).

[58] Scheffler similarly invokes emotional vulnerability as part of valuing, along with believing that something is valuable and seeing it as a source of reasons (2011: 28-31). See also Lin (2017: 373).

[59] Note that elsewhere, Street uses the term 'valuing' in a broader sense, to include a wide array of different attitudes (2016a: 177). It is an interesting exegetical question whether Street is using 'valuing' in a more restrictive sense here, to refer to one *specific* kind of complex attitude (2012: 44).

situations, but because we *value* the relationships at stake, we stay in them, working towards understanding and reconciliation.[60] Valuing, understood in this sense as a complex attitude or disposition that involves emotional commitment, is another example of a more sophisticated attitude through which we can overrule our more basic normative encounters.[61]

The upshot of this discussion is that we can exercise these more sophisticated attitudes to overrule (or strengthen) our initial encounters with normativity – what might prove to be a distinctively human ability.[62] When we use one attitude to overrule another from within the practical standpoint, we overrule a *pro tanto* reason with a stronger, more decisive reason; and when we use one attitude to strengthen another from within the practical standpoint, such as when we choose some desire, we can be said to strengthen an existing *pro tanto* reason. And all of this is to be explained in terms of *what it is* to set an internally authoritative normative direction from within the practical standpoint.

But ultimately, we are interested in understanding what our decisive stance on some normative matter is, from both a theoretical and practical perspective. So let me pose a question to the reader: can any of the attitudes we have identified – even the more sophisticated ones like choice or valuing – serve as a master attitude here, that constitutes a decisive normative stance in all different normative circumstances?

The answer to this question must be no, and that turns again on the reasons we surveyed for rejecting Attitude Monism – even for these more sophisticated faculties. For instance, I have argued that although *sometimes* our judgments can be authoritative, it is possible to overrule a particular *judgment* through *choice*, such as when we choose not to purchase what we judge to be valuable. This might lead a person to think that *choice* is the decisively

[60] Or, in a simpler example, someone who has been hurt from war might not desire to have his leg sawed off, but would still *value* that insofar as it was necessary to keep him alive (see Street 2012: 43, in an example from Korsgaard 1997: 238).
[61] In later Chapters, I employ the term 'valuing' in a broader sense (e.g. Chapters 4 and 5), encapsulating attitudes beyond just the complex one described here.
[62] This will turn on what it is like to be another animal, such as a dolphin or elephant. But as Charles Taylor suggests, "what is *distinctively* human is the power to evaluate our desires, to regard some as desirable and others [as] undesirable" (1985: 15-16, emphasis edited).

authoritative attitude in all cases, but I have argued that this is not so. Again, recall the cases of **Sunshine** and **Taste**: in moments where we have yet to make reflective choices yet, it still seems that our more basic attitudes can constitute what is good for us, as far as these experiences are concerned. And, although valuing is a sophisticated and complex attitude, it also cannot serve as the authoritative attitude in all cases: for there are many cases in which we lack the sorts of emotional vulnerability that valuing involves, yet we still make authoritative and disinterested decisions about a necessary course of action in challenging circumstances.

This means that we lack a master concept that refers to our most authoritative stance on some matter from within the practical standpoint, and which provides our decisive normative position with respect to that circumstance. But nevertheless, it seems *cogent* to ask, in any given circumstance, and on any given normative question, what our most considered or internally authoritative stance on it is. For that reason, we can put a name to the notion that can serve this conceptual function in our theorizing: I shall call our most considered or internally authoritative stance on some normative matter a *normative endorsement.*

A normative endorsement is not synonymous with the use of a specific attitude, but rather, it refers to our capacity of exercising whichever attitude is the most internally authoritative on some given normative matter. The idea of there being 'internal authoritativeness' is one that I have sought to illustrate through the various examples in which we take one thing to matter more than another, from within our experiential or deliberative perspective – from within the practical standpoint – though the notion is to some extent irreducible. Perhaps the best way to illustrate what I have in mind is through some examples of how endorsements can occur *through* different attitudes, in different circumstances:

> **Endorsement as *choice*:** Although you have judged that the item in the shop is beautiful and valuable, you *choose* not to purchase it. Here, because choice

is the most internally authoritative attitude – there is nothing from within your practical standpoint that serves to overrule the choice – your choice constitutes the normative endorsement.

Endorsement as *enjoyment*: In the first few moments of tasting the dish, you come to enjoy it greatly – even though you have not yet had the time to form a belief or come to a reflective choice. As far as the goodness of the dish is concerned, your *enjoyment* of it constitutes your normative endorsement in those few seconds.

Endorsement as *desire*: When Jane is parched and desires to drink water, she represents the drinking of water *as good* through a positive valence that is attached to it – and she lacks any further second-order desires or judgments on the matter. Here, the desire to drink water constitutes her normative endorsement of it.

Endorsement as *valuing*: When one *values* a relationship with a friend, one takes the efforts to preserve it and work through difficulties, even when one desires to the contrary. The attitude of valuing here can be understood as involving complex emotional patterns and vulnerabilities, and is therefore more involved than just a judgment or choice. But there is no further attitude within oneself that detracts from valuing the relationship, and so the valuing constitutes a normative *endorsement* of the relationship.

What these examples are meant to illustrate is that we can make normative endorsements – i.e. internally authoritative normative verdicts – in a wide range of circumstances, in which *different* attitudes can serve as the crowning normative act. The reason it is useful to formalize the notion of an endorsement, however, is that it is the authoritativeness of an attitude that provides a presumptively decisive normative verdict on some matter, from within the practical standpoint. Endorsements are what either allow us to sift through and weigh our preliminary *pro tanto* reasons for action, or in cases where we are unable to make those more sophisticated reflections, will prove to *just be* one of

our more basic attitudes. This is why animals can also have normative endorsements, but through their more basic attitudes, rather than through higher-order reflection: for animals, their basic enjoyments and aversions and desires will in most cases constitute their normative endorsements, i.e. the most authoritative direction setting from within their practical standpoints.[63]

Our distinctively human ability to make endorsements through our more reflective capacities is what makes us not practical automatons, or beholden to our more specific attitudes. And this is an important conceptual point, against certain theorists who claim that our normative reasons are grounded in some single attitude. For instance, imagine some Humean who claims that all normative reasons are grounded in desires. This can sound like an antirealist claim, and it would imply that our decisive reasons are also grounded in our desires. But then imagine that someone perfectly understands this claim, and professes the opinion that she *doesn't* think that she has decisive reasons to follow her desires; she prefers instead to follow her judgments, or her choices, or so forth. If the Humean insists that she nevertheless *does* have decisive reason to follow her desires – that is the Humean claim – then the Humean will start to resemble a dogmatic realist: one who insists upon some normative truth as if it were a *brute* truth, against the conclusions that an agent comes to from within her practical standpoint. It is again, for such a reason, that Humean claims of this kind, alongside other forms of Attitude Monism, are not genuinely *antirealist* positions that avoid brute normative claims and which respect the conclusions of the practical standpoint. Our account of Attitude Pluralism, which puts normative endorsement at the center, better permits us to take the reins of our lives, for human normative endorsements can attach to higher forms of reflection that most represent 'what we stand behind.'

[63] Although I have focused on cases of internal direction-setting in a positive sense, we can also speak of normative *rejection*, which is when we take an internally authoritative stance on the *badness* of something. For instance, if an individual finds some music to be bad – whether that is through displeasure, judgment, emotion, or otherwise – her most *internally authoritative* stance on the music, as a bad piece of music, would constitute her normative rejection of it.

The picture I have provided means that we are normatively free, but that does not mean that we choose all of normativity.[64] Much of the normative world is presented to us in a passive, involuntary fashion – from within the practical standpoint, the irreducible goodness of enjoyment and love, as well as the irreducible badness of pain, do not feel like things we have *chosen* as good or bad, but rather things that present themselves to us as bad. Much of normative life is like this – we stumble across things that appear normatively laden, where the phenomenology of that does not involve much active effort on our own parts.[65] But ultimately, we are not beholden to those involuntary normative responses: for higher forms of endorsement can happen through our more agential capacities, of valuing and judgment and *choice*, and that is what makes us normatively free as humans.

One benefit to having the notion of normative endorsement is that, in any given circumstance, we need not primarily ask the question of which attitude is being exercised – which is the question that the literature has focused on – but rather, we can focus on the normative answers that individuals would provide to complex normative questions, answers which usually reflect their authoritative normative positions. When someone is in the sunshine, and we ask them, "How does it feel?", their simple enjoyment will speak for their endorsement. When we ask the question "I know that this experience brings you pleasure, but *how much* does pleasure matter to you, among other normative goods?", we will again have their normative endorsement, whether that is through judgment or choice.[66] We can furthermore ask things like "What do you care about most?" or "How much do you holistically like the state of life you are in right now?" – and from both a practical and theoretical perspective, it will be their normative endorsement on the question that matters.

---

[64] We are normatively *free*, in the sense that our reasons are not beholden to any specific attitude – we can override them from within the practical standpoint ultimately through choice. But bear in mind we have not yet seen the constitutive arguments of Chapters 3 and 5 that will amend our reasons.

[65] Again, this is why *constructivism* is not a perfect term for our project, for it connotes too much construction on the agent's part (§1.4).

[66] This is why some individual *could* endorse hedonism, but also why we are not all *beholden* to hedonism, as a truth about our normative reasons or prudential goodness.

Importantly, it will be an individual's *all-things-considered* endorsement that will constitute her presumptively decisive reasons for action. A normative endorsement is an individual's most internally authoritative stance on some given normative matter; but an *all-things-considered* endorsement is the most internally authoritative answer she provides to the general question of 'what do you take yourself to have most reason to do in this situation?' I say that this explains her *presumptively* decisive reasons for action, because we need to qualify that we have not yet considered the constitutivist arguments I will provide in Chapters 3 and 5 that will provide sources of further normative correction for her reasons. But when an individual looks at a situation with its holistic features in place, and comes to an overall normative stance on *that*, she will have rendered an all-things-considered endorsement, and constituted her presumptively decisive reasons.[67]

As a final point about normative endorsements, it is important to note that they are by no means limited to considerations about *prudential goodness*, also known as what is good *for* an individual, which refers to something like egoistic or personal benefit.[68] Although it is true that I have been discussing normativity in a way that does not yet make a necessary reference to moral considerations – a task we take up in Chapter 5 – this does not mean that our normative endorsements can only target personal or egoistic benefit. For there are many things that we can endorse *as good*, which extend beyond ourselves: these include certain selfless ways of life, the flourishing of others, long-term projects, musical or artistic pursuits, and so on. After all, recall that the point of a normative endorsement is to provide ourselves an internally authoritative source of normative reasons. And so, Williams was right to note that "there is of course no supposition that the desires or projects of an agent have to be egoistic; he will, one hopes, have non-egoistic projects of various kinds, and these equally

[67] We should leave conceptual space for the possibility that sometimes our endorsements are indeterminate, and also that we may not always be aware of our normative endorsements, or our other attitudes – i.e. that we are not perfectly internally luminous. For instance, Williams writes that someone could be "ignorant" of some attitude they had (1981: 103). See also Korsgaard (1996: 92), as well as my discussion in §5.3.

[68] See Parfit (2011a: 39) and Crisp (2021) for discussion of this widespread notion.

can provide internal reasons for action" (1981: 105).[69] To describe our overall set of endorsements, it can therefore be useful to speak of what is good *from* the point of view of someone – rather than just what is good *for* someone – so as to capture the grounds of all an individual's (pre-moral) reasons, which need not be directed solely at one's own egoistic interest.[70]

This is why the *dualism of practical reason* – the purported tension between two rational principles, one solely concerning our egoistic benefit, and the other concerning universal benevolence – is not a necessary feature of our mind-given theory of normativity. Sidgwick called this dualism "the profoundest problem of Ethics" (1907: 384n4), but it only becomes such a problem if we take on certain substantive normative presuppositions – presuppositions that are not necessarily natural to the ordinary endorsements of all human beings, and which are not legislated by the nature of our practical standpoints. This does not mean that the categories of self-interest and universal benevolence are not cogent as categories – it just means that the tension between them only turns into the "profoundest" problem of Ethics, insofar as we choose to regard it that way.

2.4 Theoretical Upshots

Let me now describe some of the theoretical upshots of this picture. First, I want to show how this provides a more plausible starting point for a mind-given theory of normativity, against certain objections that Parfit (2011) and Scanlon (2014) have leveled at antirealist or subjectivist theories; and second, I want to explain why our view on the relationship between desire and choice is different from certain remarks found in Korsgaard (1996).

---

[69] Indeed, for those who have meditated upon Scheffler's insightful thought experiment – of how our priorities would change, if we knew that the human species would go extinct a few decades after our own deaths (2013) – it will become clear that most of us care much more than about ourselves.

[70] Cf. Korsgaard's usage of 'good for' to capture what I describe here as 'good from' (2018); see also Hurka (1987: 73) and Fritzson (2016).

First, we have opened up our mind-given theory to a wide range of normative inputs – and this should help to assuage concerns about views that are focused only on desire. For instance, both Parfit and Scanlon have protested that certain subjectivist or internalist views emphasize desires at the expense of enjoyments. Parfit writes that "many people fail to distinguish between hedonic *likings* or *dislikings* and such meta-hedonic *desires*" (2011a: 54, emphasis added) and insists that "it is our hedonic likings and dislikings, not our meta-hedonic desires, that *make* these conscious states good or bad" (2011a: 3, emphasis added). Scanlon makes a similar point against Mark Schroeder:

> Much depends on his claim that facts about desires, *rather than, say, about what a person would enjoy*, are the best explanation of the reasons that one person (who likes dancing) has to go to a party where there will be dancing (as compared to the reasons of another person, who does not like to dance). This claim does not strike me as plausible. (Scanlon 2014: 6-7, emphasis added)

But as I have argued, there is no reason to think that our enjoyments or likings cannot serve as the ground of our *pro tanto* reasons.[71] Because our enjoyments and likings involve experiencing an irreducible goodness from within the practical standpoint, they do always constitute *pro tanto* reasons, and when they are either left alone or further strengthened by judgment or choice, they can even constitute our presumptively decisive reasons for action.[72]

Enjoyments and likings therefore constitute an important way in which the world becomes a normative place for us – and indeed, Parfit makes some remarks here that make him sound quite like a mind-given theorist. For when he

[71] See Sobel for a similar point: "subjectivists should claim that intrinsic favoring and disfavoring attitudes ground reasons when they are accurately informed about what their object is like" (2016: 281). Of course, Parfit and Scanlon might still insist that desires themselves *never* ground reasons, and I have explained why our mind-given theory can disagree with this. But the point is that Parfit and Scanlon cannot object that our account *excludes* enjoyments as sources of reasons.

[72] Parfit and Scanlon often take these points against desire-based views to count in favor of their own forms of realism; but one of my points here is that the rejection of exclusively desire-based views does not yet constitute a reason to adopt realism, for there are other positions available, which are not exclusively desire-based.

encounters the question of whether we have reasons to like or dislike things at all, he makes our *likings* and *dislikings* fundamental: "*hedonic likings* or *dislikings* cannot be rational or irrational, since we have no reasons to like or dislike these sensations" (2011a: 2-3). But an incredibly salient point here, which should be apparent to readers of these lines, is that likings and dislikings are *attitudes* of ours: attitudes which Parfit says are what "*make* these conscious states good or bad" (2011a: 3, emphasis added), and where there are no prior reasons to *have* these attitudes. This, to me, seems like a key vindication of an antirealist or mind-given insight, and it is one that Hume articulated long ago:

> If you push your enquiries farther, and desire a reason, why he hates pain, it is impossible he can ever give any… Something must be desirable on its own account, and because of its immediate accord or agreement with human sentiment and affection (Hume 1751: 87).

Hume's point, like Parfit's, is that there is ultimately no prior reason to have some particular liking or disliking – rather, it is our basic sentiments and affections that ground our reasons. Here, it seems that we find a kernel of antirealism in Parfit.[73]

Second, I want to emphasize how the positive picture we have advanced stands in contrast to certain remarks from Korsgaard, who can write at times as if there is *no* normativity at all, prior to the act of reflective choice. For instance, she writes fairly explicitly that desires do not provide a person with reasons, such that if someone does not ever make a reflective choice, she will have *no* reasons at all:

> The normative sceptic has no reason for doing anything. He does have some desires and impulses, of course – those are to some extent supplied by nature –

[73] Street insightfully latches onto this point, writing that "Parfit is explicitly acknowledging that what attitude-dependent theorists argue is *always* going on in cases of value does at least *sometimes* go on, namely that value or disvalue is conferred upon the world (in this case, upon our own sensations) by evaluative attitudes which in and of themselves are ultimately neither rational nor irrational" (2009a: 283). See also Frankfurt (2004: 26).

> but he has no reason for acting on one rather than another… the practical normative sceptic does not really have any *ends*, since his desires do not provide him with reasons. (Korsgaard 1996: 163)

Or, as Korsgaard emphasizes elsewhere, "the reflective structure of human consciousness… is the source of normativity" (1996: 103-104) – where this is pointing to our capacity for reflective *choice*, rather than other faculties like desire. This is a way of reading Korsgaard as an Attitude Monist about *choice*. And such a thought, which emphasizes normativity from the *will*, may be traced back to suggestions from Kant.[74]

But if we have read Korsgaard correctly here, then we can find certain counterintuitive implications from such a position.[75] One counterintuitive implication is that, when humans are in tired or exhausted moments – moments in which they struggle to really make a reflective *choice* – they have no reasons for action. Suppose that in **Parched**, Jane has a strong desire to drink water, but her mental faculties are too weakened in order to reflect upon the desire from a standpoint of choice: as I have argued, it seems like the desire alone is sufficient for grounding a reason. Another counterintuitive implication of this position

---

[74] For instance, Kant discusses the idea of a will that sets aside *all* desires and sensuous stimulations:

> That is why a human being presumes for himself a will that lets nothing belonging merely to his desires and inclinations be put on its account, and on the contrary thinks possible – indeed even necessary – through himself actions that can be done only by setting aside all desires and sensuous stimulations. (Kant G. 4: 457)

This would be a will that can make pure choices, while letting "*nothing*" in the realm of desires and inclinations "be put on its account." Although Kant does not explicitly claim here that the will is an *exhaustive* source of normativity – with desires having *no* normativity to them at all – his remarks about the will do make it easy to see how one could find such a suggestion in him. Indeed, perhaps in recognizing the challenging nature of this view, Kant says that a satisfying explanation for the possibility of such a will can "never be understood by any human reason" (G. 4: 460).

[75] This reading of Korsgaard might be softened in two ways: first, it might be the case that Korsgaard does not use the concept of a *pro tanto* normative reason, and therefore sees 'reason for action' as implying what others would describe as a *decisive* reason (compare 1996: 163 and 2009: 52). Second, Korsgaard sometimes writes as if neglecting to think too hard about a desire just does constitute a reflective endorsement: "For in one sense no human action can happen without reflective endorsement. When people skip reflection or stop too soon, that is a kind of endorsement, for it implies that the work of reflection is done" (1996: 161; see also 127 and 140). That would, of course, be a significant claim.

would be that there is no normativity in the lives of many animals, who cannot make reflective choices: for it is only reflective *choice* that generates normativity, then the pleasures and pains of these animals would not be sources of normativity for them. Korsgaard later makes admirably clear that she does think there is normativity in the lives of non-human animals (2018), but on the basis of her earlier remarks from *Sources* (1996), it is less clear that the view makes space for such normativity in the lives of animals.

Generally speaking, it seems important to vindicate that there is normativity prior to reflective choice. Not only does this makes sense of a variety of our experiences that seem *normative*, even prior to reflection – when we walk into the sunshine, when we are too hungry to think, when we are in too much pain to deliberate – but this is also a crucial point for choice itself, for pre-reflective normativity is what gives the will *something to choose.* If there is no normativity prior to choice, then choice becomes something numb, something without a compass: it only points blindly or arbitrarily in different directions. It is the fact that our basic enjoyments and aversions and desires provide us with initial normativity that choice *has something to choose at all* – and that is a further reason why it is a more attractive view to think that choice can *overrule* or *strengthen* our pre-existing pro tanto reasons, but not serve as the sole source of our reasons.

2.5 Normative Lens

I have spent most of this Chapter discussing the nature of our normative attitudes, arguing against Attitude Monism in favor of a distinctive account of Attitude Pluralism. In this final section, let me make a broader observation, which is to point out that an analysis of our practical standpoints is not exhausted in a discussion of our normative attitudes: for there is a background dimension to our practical standpoints that *precedes* considerations of the specific attitudes that we can experience.

Let me call a *normative lens* the background conditions that structure *how* and *when* particular attitudes appear to us. This is a notion that I think mind-given theorists should pay more attention to. For instance, the very fact that we see certain objects or persons in the world as presenting us a certain set of opportunities for action is, itself, a feature of our practical standpoints that could have been different. The fact that one person might be struck by a desire to *hug* someone when seeing them, and another person to *bow* when seeing them, represents a difference in the normative lens that two people occupy. The fact that a chair is even something that presents itself to us as *to-be-sat-on*, which might be described as a minimal layer of normativity, is something that is peculiar to the *human* practical standpoint; dragonflies and cows do not see the chair in that way. (The fact that pigs see the mud as something *to-be-rolled-in*, illustrates the difference between their normative lens and ours as well.) The possible examples and illustrations here are pervasive enough that it is sensible just to summarize the basic point as follows: there are many background conditions that go into the formation of our more specific desires, judgments, and choices, conditions which are importantly a part of the practical standpoint that we bring into the world. These conditions color the very nature of what sorts of normative attitudes we are likely to experience, what we regard as opportunities for action, and even what concepts we use to render our own experiences intelligible to ourselves.[76]

This point is of significance to an account of normativity, for if we want to direct our lives in a normative fashion, then we must pay attention not only to our particular attitudes themselves, but also to the internal patterns, tendencies, and concepts that *condition* how those attitudes arise – and which are not always *necessary* features of our practical standpoint, but contingent ones. The normative lens that we come to have may be a product of biological,

[76] This is especially true with the use of *thicker* normative concepts, which might vary significantly between different times and cultures. For instance, if one set of people routinely uses the normative notion of *repentance*, while another set of people routinely uses the normative notion of *filial piety*, these can create different patterns that structure their normative experience of more specific desires, thoughts, and feelings.

environmental, and social factors – and this allows us to possibly adjust our normative lens over time.

Let me give you a concrete example to elucidate the notion of a normative lens. Suppose that someone is misogynistic, and has work to do on better understanding the dignity and fundamental equality of women. This person cannot *fully* accomplish this task simply by deciding what to do with the particular desires, thoughts, or judgments that come to mind. This person also needs to take a long look at the background conditions within the practical standpoint that give rise to those specific attitudes: the *normative* connotations that might come along with that person's very *concept* of a woman; the habitual or ingrained *tendencies* of thought or feeling that lead to the appearance of specific normative attitudes; and what are the fullest horizons of action, perhaps previously unconsidered, that are possible.

Another example has to do with music: people who have acquired different background conditions in their practical standpoint have different *normative* experiences of music. Acquaintance with certain pieces, as well as familiarity with certain musical concepts, can change a person's perception of the *beauty* in a piece of music. It is not enough for the philosopher to point out that one person experiences a judgment of beauty in a piece, while another person does not, and to conclude that the normative difference here lies in the fact that one person has a normative attitude, while the other does not. There is a deeper and relevant difference between the two people, which is that they occupy different normative lenses that *structure* what is possible in their normative encounters with music, and what judgments they are disposed to come to, or are even capable of. And if we are interested in the realization of certain normative goods, we must pay attention to questions about how we can come to influence the nature of our normative lens.[77]

[77] My notion of a normative lens has some connection to ideas in phenomenology, such as when Dreyfus interprets Heidegger as discussing "a background understanding of what counts as things, what counts as human beings, and ultimately what counts as real, on the basis of which we can direct our actions towards things and people," i.e. a "clearing" (1995: 27). The notion can also be compared with the idea of *affordances* in psychology (see Gibson, *The Ecological*

Conclusion

This Chapter focused on developing an account of how the mind renders the world into a normative place. I argued against the widespread thesis of Attitude Monism, in favor of a distinctive account of Attitude Pluralism. According to the view I have advanced, our most authoritative stance on some given normative matter is represented by our *normative endorsement* on that matter – and our presumptively decisive reasons for action are constituted by our *all-things-considered normative endorsements*. This has hopefully re-oriented our focus away from the narrow question of *which* specific attitude matters for normativity, to the broader point that human normative life can be rich and multi-faceted. Finally, I have argued that our experience of normativity is deeply colored by the normative lens that we occupy.

In the next Chapter, I build upon this account in arguing that certain features in our practical standpoint commit us to deeper normative principles that create a structure of normative objectivity in the pre-moral domain.

Approach to Visual Perception), as well as with recent work in conceptual engineering and social metaphysics (see Barnes 2017: 2422, discussing Haslanger).

3 The Structure of Normativity

3.1 Introduction

One of the main challenges to a mind-given theory of normativity is that it seems to provide the wrong verdicts on substantive normative matters, such as what we have most reason to do. This is because, despite everything we have said about normative endorsements in the previous Chapter, our normative endorsements can still go in any direction. It is possible that someone endorses something that is detrimental to their own health or happiness; and it is also possible that someone endorses something that is detrimental to others. Are we really to say that these people therefore have decisive normative reasons to follow through on such endorsements? Because a mind-given or antirealist theory claims that our reasons are grounded in our endorsements, this seems to be the result.

Here is how Ruth Chang articulates the general form of this challenge, describing what we have called a mind-given or antirealist theory under the heading of 'source internalism':

> Source internalism fails to guarantee the right substantive results about what reasons we have because the constraints it puts on desires are purely formal in nature. No formal constraint, however intricate, can *guarantee* the intuitively right answer as to what reasons we have. (Chang 2013: 173)

Here is how Street articulates the same kind of challenge to antirealism:

> It's easy to see how one might argue *that there are external reasons*. To undertake this, the general method is to imagine cases in which it does *not* follow from someone's subjective motivational set that he or she has normative reason to act in a certain way, and then call upon people's intuitions to the effect that nevertheless the agent *does* have reason (in the standard normative sense) to act in this way. (Street 2017: 133)

There are many examples in the literature that motivate this sort of objection. For instance, consider Rawls' case of someone who only takes pleasure in counting blades of grass:

> Thus imagine someone whose only pleasure is to count blades of grass in various geometrically shaped areas such as park squares and well-trimmed lawns. He is otherwise intelligent and actually possesses unusual skills, since he manages to survive by solving difficult mathematical problems for a fee… Naturally we would be surprised that such a person should exist. (Rawls 1971: 379)

Numerous commentators have maintained that such a case suggests that antirealism is mistaken, on the basis of an intuition that such an 'intelligent' person has more normative reason to do something else besides counting grass – for instance, to spend his time studying pure mathematics instead. The problem is that antirealism seems to fail to deliver this result.[78] Other cases in the literature involve Anscombe's person who desires a saucer of mud (1957: 70), Gibbard's anorexic who wishes to starve herself (1990: 171), Wolf's man who lives to make handwritten copies of *War and Peace*, and the woman whose "world revolves around her love for her pet goldfish" (2010: 16). The point of these examples is to show that our actual endorsements can lead us in some odd and objectionable directions.

More deeply, even when we set aside these cases, we can see that there is a more general need for normative objectivity. If all a theory told us is that we should follow our endorsements, it wouldn't be telling us much that is ultimately interesting or helpful: for much of everyday life is about seeking to figure out what we *really* have most reason to do, even when we are already aware of where our attitudes are seeming to point us. When you are deciding what job to take up, where to move, or who to spend your life with, you want to know more than

[78] This now famous case is discussed in Heathwood (2005: 499), Street (2009a: 274), Bruckner (2016: 3), and Crisp (2021: 4.2), among other places. The case can be traced back to a suggestion in Stace of someone who counted "exactly 1,422,769 blades of grass on the lawn" (1944: 238).

just what you already appear to think – you are searching for some *right* answer that extends beyond your current attitudes.

The realist's response to this situation is to assert that there are certain brute and mind-independent normative facts that determine what we have most reason to do, and which serve as correctives towards our actual attitudes. For instance, the realist might assert that the grass counter *just is* mistaken about what he has most reason to do, or about what has most intrinsic value – even if the realist cannot ultimately offer some more fundamental explanation for why her claim is true. As I stated earlier, I do not want to claim in this dissertation that realism is false – but because it is a dogmatic position, as well as one that gives rise to other metaphysical, epistemological, and explanatory objections, it behooves us to see whether there are other ways of arriving at objective normative correction (§1.2).

Second, some antirealists, by contrast, have suggested that we bite the bullet: that what you have most reason just is what you take yourself to have most reason to do, without further sources of normative correction beyond that. This would be a dismal result if it were true, for there certainly *seems* to be more to normativity than that.

A third and influential approach has been to claim that our normative reasons, or prudential goodness, are in some sense linked to our normative attitudes, but that this is only so after we *idealize* upon our attitudes. Such views are widespread, and we can take Railton and Smith as two paradigmatic examples:

> Let's say that an individual's intrinsic good consists in attainment of what he would in *idealized* circumstances want to want for its own sake… were he to assume the place of his actual self. (2003: 54, emphasis edited)

> The Dispositional Theory… tells us that facts about final goodness, as indexed to an agent, are fixed by the facts about the final desires of that agent's *ideal* counterpart. (Smith 2013: 18, emphasis added)

The thought here is that, if we propose certain requirements that would correct for our starting attitudes – requirements of full information, internal coherence, or the like – then we can achieve the right kinds of substantive verdicts.[79] Seen in the light of a quest for the right substantive verdicts, such idealization can seem like an intuitively attractive approach. But all idealizing views face the following objection: where does the *normativity* of the proposed idealizing requirement come from? Suppose that some agent fully understands the proposed requirements, or understands that a certain conception of his ideal self wants something different than himself, but *denies* that he has any reason to follow this idealized advice. We can then ask, with Roger Crisp, the following:

> If we allow that I should φ only in so far as a certain procedure has as its upshot that I should, do I have reason in the first place to act on the basis of the deliverances of the relevant procedure?... If the answer is affirmative, then we have an INE [intrinsically normative entity]… If the answer is negative, then we have nihilism. Why should I act on the deliverances of some procedure if I have no reason in the first place to do so? (Crisp 2006: 53)

This objection asks us to consider that many idealizing views, though they attempt to draw some connection between our normative reasons and our attitudes, are at bottom *realist* views: for if the idealizing procedure is asserted as a normative requirement, without some further explanation, then it will just turn out to be a brute and mind-independent normative fact. It might be presented as a *procedural* (or *formal*) rather than *substantive* requirement – one that tells us a general constraint upon our reasons, rather than to directly hold one reason rather than another – but at the end of the day, the alleged *normativity* of the requirement will beckon towards the same question. That is the question of where the normative fact comes from – the normative fact saying that we ought

[79] Idealizing views can also be found suggested in Sidgwick (1907: 111-112), Rawls (1971: 407-424), Brandt (1979: 10, 113, 219), Hare (1981: 101-105), Griffin (1986: 11-17), Kagan (1989: 283-291), and Lewis (1989: 113-137), among others.

to follow this particular procedure, which the individual might otherwise claim to reject as normative for her.[80]

As I shall suggest, a promising way out from this situation is to explore whether there are certain principles that humans are either *in fact* or *constitutively* committed to that can serve as sources of normative correction for their starting attitudes. If some principle is something that a person is, *as a matter of fact*, committed to, then there will be no further need to justify the principle to them – it will be a principle that emerges from her own practical standpoint. If some principle is something that a person is *constitutively* committed to, then it will be something that is true for them *in virtue of* an existing commitment of theirs – it will just be a matter of exposing to them that there is some principle that they already maintain as true. So, if we can find certain normative principles that can be justified on either of these grounds – and which can offer us a way of correcting our starting attitudes in a way that achieves intuitive verdicts in a wide range of cases, and which can explain our pre-theoretic sense that there is a hierarchy to normativity that extends beyond our own current attitudes – then we can offer a path forward that is neither explicitly realist, substantively dismal, nor covertly realist.

My central aim in this Chapter is to outline a view that accomplishes that task in the pre-moral domain. In the following section, I will argue that all humans, either as a matter of empirical or constitutivist fact, are committed to the normative principle that they have most reason to pursue their strongest *counterfactual* endorsements (§3.2). I argue that this means we are all committed to an objective hierarchy of mind-world potentialities – objective in the sense that this normative hierarchy is independent of our current attitudes about what we have most reason to do – and therefore that the mind-given or antirealist theorist can promise to achieve a significant degree of normative objectivity. After outlining and arguing for that central principle, I then discuss some further

[80] Although my aim here will not be to parse through the plethora of idealizing views that exist in the literature, it seems clear that many of them do not offer antirealist or constitutivist justifications for the normativity of their requirements.

ways in which we can make mistakes with reference to our normative endorsements (§3.3). I then discuss how these principles apply to substantive matters, including the grass counter (§3.4).

The biggest caveat I must make at the outset is that, in this Chapter, I intend to discuss these questions in a way that is independent of, or prior to, *moral* considerations. This is not because morality is not important – indeed, it is the most important topic of this dissertation – but because it is possible to ask the question of what we have most normative reason to do, in a way that does not yet make a *necessary* reference to morality. For instance, many authors ask the question of what is good for someone, in a prudential sense, without yet looking at the question of what is moral; and consider that Aquinas discussed what commentators describe as "pre-moral" principles, as part of his conception of natural law.[81] Although I do not have further space to discuss this conceptual assumption here, I shall take it as granted for the sake of our discussion. In Chapter 5, I will then discuss how moral considerations further inform, enrich, and constrain our normative reasons – which is what will bring our account of normativity into full fruition.

3.2 The Standard of Possibility

What I shall now argue is that, either as a matter of empirical or constitutive fact, human beings are committed to the principle that they have most reason to pursue their strongest counterfactual endorsements. I call this the **Standard of Possibility**: it holds that, whenever we are deliberating about what to do, we commit ourselves to the idea that, among the possible courses of action that are available, it is the one that we *would* endorse the most, upon realizing that path, that provides our most decisive normative reasons. Because I shall also argue that many of the potentialities surrounding us have an objective dimension to

[81] See Aquinas, *Summa Theologica* I-II q. 94, as well as John Finnis: "Principles of right and wrong, too, are derived from these first, *pre-moral* principles of practical reasonableness" (2011: 34, emphasis added; see also 32).

them – the degree to which we would counterfactually endorse something is not dependent upon our current attitudes or predictions – this principle, if given a successful defense, will serve to connect us to an objective hierarchy of normative possibilities. Let me first explicate what I have in mind with this principle, at the level of content – before I then move into the two modes of empirical and constitutivist justification for it.

The principle is quite simple in its essence. Every time you are deliberating about what to do, there is always a range of possible courses of action that are available – this is simply a metaphysical observation about the fact that possibilities always surround a person. The particular contents of those possibilities will vary depending on a person's contingent location in time and space, but the fact that there always *are* possibilities is a universal fact: *you* might be deciding whether to order the pad thai or green curry tonight, while Lincoln might have been deciding whether to run for President or not, while Genghis Khan might have been deciding which new territory to invade. There may be differing extents to which some possibilities are more achievable than others, though this qualification about achievability will only have limited relevance for our core arguments. The observation here is that *all* points in time and space involve multiple possibilities for action.

Next, notice that there are endorsements that we *could* be making from the standpoint of having taken each of these potential paths. The endorsement that we would make, upon attaining the aim of some possible action, is what I shall call a *counterfactual endorsement*: the endorsement that is taken from that potential standpoint reached by the action.[82] So, for instance, there is a counterfactual endorsement as to how much you would enjoy the pad thai, and another one as to how much you would enjoy the green curry – and these counterfactual endorsements can be thought about, even before we actually try any of them. This is because our counterfactual endorsements often have a

[82] I am using the term 'counterfactual' in a somewhat unorthodox manner, when it often refers to what *could* have been, in a retroactive manner. But I have chosen it precisely for this reason: more than the terms 'possible or 'potential,' it signals the sense in which there is comparative pressure – and potential regret – to the fact that there are different paths available.

determinate quality to them: although this is not the case for *all* of our normative reactions, the combination of some set of internal normative dispositions, along with some state of affairs in the world, often produces a determinate normative reaction in result. What will your normative reaction be if you put your hand in the fire, as opposed to a cool spring of water? What will your normative reaction be if you try out a vocation that you have no interest or passion for, versus another one that matches your personality and cares? The answers to these questions involve determinate answers about the *mind-world potentialities* that surround a particular human: given that we are a certain way, with certain biological, psychological, socially-informed practical standpoints and dispositions, there arise also certain facts concerning the nature of our counterfactual endorsements.[83]

These counterfactual endorsements often come in different degrees of *strength*. There are some things that we can like more than others, or some things we can take to be more worthwhile or good than others. The question of what it is to holistically endorse one state of affairs more than another is one that may have some irreducibility to it: we all know what it is like to endorse one experience more than another – for instance, from when we reflect upon our past memories, good and bad – though it can still be hard to explain what this notion of comparative strength amounts to in explicit terms.[84] Not all comparisons between endorsements have a determinate quality to them – some might be on a par,[85] or otherwise indeterminate, especially if the counterfactual endorsements are themselves indeterminate. But a wide swath of our comparative

[83] It is possible that, at some counterfactual standpoint, our basic dispositions would have undergone some change – but this is intended to be a feature of the view, including if we undergo transformational change (see Paul 2014).
[84] *Comparisons* between the strength of counterfactual endorsements can be understood either in absolute terms, or in terms of being made from some heuristic standpoint. The point is that we all know what it is like to take ourselves to more strongly endorse one thing rather than another – for instance, endorsing the pad thai more than the green curry – though it can be hard to define this point in further analytic terms. That difficulty does not mean degrees of strength to our endorsements are not cogent. See also Chang (2016) for discussion of the nature of a comparative fact.
[85] See Chang (2002).

counterfactual endorsements are determinate, and that is the relevant point to secure here.

The Standard of Possibility maintains that we have more reason to pursue our stronger counterfactual endorsements – and, by implication, that we have most reason to pursue our strongest counterfactual endorsements. This should be understood in terms of *holistic* counterfactual endorsement, from the standpoint achieved by some possible action – i.e. with all normative features of the situation in mind. It should also be understood as the counterfactual endorsement, or set of counterfactual endorsements, that are made in relation to the casual effects of the action in question.[86] So, to take a general example: suppose that there are two doors in front of you – you do not know what is behind each door – but as a matter of fact, entering the first one entails smelling some aroma that you would like, while entering the second one entails experiencing some excruciatingly painful electronic shocks. Suppose that you must choose one. Your stronger holistic counterfactual endorsement lies in entering the first door, rather than the second – and the point of the Standard of Possibility is to establish that you have more reason, on the basis of that stronger holistic counterfactual endorsement, to enter the first door, rather than the second. But notice that this point holds even though you do not *know* which door contains which result – for the point of this principle is that you are committed to ascribing more reason-giving status to entering the door that involves the stronger counterfactual endorsement, *whichever* that door turns out to be. Your commitment to the Standard of Possibility, coupled with the empirical facts, combine to give you stronger reason to enter the first door, even if you don't know that yet, as an epistemological matter. From this bare-bones example, a general point emerges, which applies to a wide range of cases, about what

[86] For instance, the holistic counterfactual endorsement surrounding the decision to eat green curry will involve the standpoint of having tasted it, coupled with the other normative upshots of having eaten the green curry, for instance on health or other dynamics at the dinner. The set of holistic counterfactual endorsements concerning a vocational choice, which is a much larger decision, will involve different endorsements that arise from the dimensions of our future practical standpoints, causally related to our vocational choice. The counterfactual endorsement should therefore be understood as being made from a standpoint after the full effects of an action can be taken into account.

vocation to pursue, who to spend time with, what to eat for dinner, and so on. Indeed, the Standard of Possibility even applies to our having most reason to *see* the world in certain ways, i.e. to adopt certain normative lens and habits of practice, that are most internally endorsable from those standpoints.

The idea that counterfactual acquaintance with some possible outcome holds an important relation to its goodness is not a new one. For instance, Plato can be read as making such a suggestion towards the end of the *Republic*, when he suggests that the person with the best life has the most robust counterfactual acquaintance with the lesser forms of life, while those with the lesser forms of life do not have the same counterfactual acquaintance with the best form of life.[87] Mill also appeals to the idea of a comparative fact that comes from acquaintance with different options:

> Of two pleasures, if there be one to which all or almost all who *have experience of both* give a decided preference, irrespective of any feeling of moral obligation to prefer it, that is the more desirable pleasure. (Mill 1863 [2015]: 122, emphasis added)

Though there are differences between Mill's proposal and ours, there is a general suggestion here about counterfactual acquaintance being important.[88] Sidgwick makes another reminiscent suggestion:

---

[87] Plato imagines a disagreement between three kinds of people about the best life: the life of learning, the life of victory, or the life of wealth. Plato's point is that only the person with experience *of all three* can be the best judge of which is the highest kind of life, and that is why he stresses that only the person who values the life of learning has counterfactual acquaintance with the other two kinds of lives, but not vice versa (see *Republic* 582). Of course, Plato is most plausibly read as making an *epistemic* point, not a *grounding* one. But we might take the following suggestion from him: though the wealth-lover does not realize it, he has a stronger *counterfactual endorsement* towards the life of learning – and the reason he doesn't realize it is that he has no counterfactual acquaintance with it.

[88] Mill's proposal differs from ours in more than one way: ours does not contain a majoritarian point about "all or almost all" agreeing on the decided preference; Mill's point is about the desirability of two *pleasures*; and it is also an open question whether Mill is making an *epistemological* or *grounding* point in his statement, whereas we are interested in something closer to the latter, in trying to use counterfactual endorsements to *explain* the greater reason-giving status of one thing over another.

> For it is not even sufficient to say that my Good on the whole is what I should actually desire and seek if all the consequences of seeking it could be foreknown and adequately realised by me in imagination at the time of making my choice… it would have to be said that a man's future good on the whole is what he would now desire and seek on the whole if *all the consequences of all the different lines of conduct open to him were accurately foreseen and adequately realised in imagination* at the present point in time. (Sidgwick 1907: 111-112, emphasis added)

Sidgwick's proposal appeals to the idea that one's highest good is based in what one would seek, after a full acquaintance with the implications of all the "different lines of conduct open to him" – a point that resembles our suggestion about counterfactual endorsement.[89] The point is that some general idea about counterfactual acquaintance is a familiar one from these and other thinkers.[90]

But it should be no surprise that neither Plato, Mill nor Sidgwick – along with many contemporary proponents of idealizing requirements – attempt to justify such principles in constitutivist terms. This brings us to the task of justification, as we want to avoid the appeal to brute normative requirements. Imagine that someone claims to understand the Standard of Possibility, but denies its normative force upon her: we imagine her to say something like, "I know that I would holistically endorse the green curry more than the pad thai, but I still take myself to have more reason to order the pad thai." The question now is: how can we show, on mind-given presuppositions, that such a person is making a mistake?

---

[89] There are differences between Sidgwick's suggestion and ours: notably, Sidgwick seems to tell us to assess these different lines of conduct *from the present vantage point*, whereas our proposal is about comparing some more determinate degree of strength in our counterfactual endorsements, *from the vantage point of those counterfactual standpoints.*

[90] For instance, the idea of counterfactual acquaintance – though it is not always put in those terms – is found in Rawls (1971 [1999]: 366), Lewis (1989: 121), under the name of *imaginative* acquaintance, and Railton (1986: 16). However, what these three distinguished authors emphasize less is the idea of counterfactual *comparison* among a range of possible outcomes, as is more explicit in Mill and Sidgwick. Moreover, none of these authors discuss constitutivist justification, as we will.

Let me now offer two independent arguments for this conclusion: the first essentially claims, upon empirical grounds, that no person could genuinely make such a claim; and the second essentially claims that such a person is engaging in a self-contradiction, i.e. making a constitutive mistake in what she is saying.

According to the first argument, which has at its core an empirical foundation, all human agents just *are* committed to the idea that they have most reason to pursue their strongest counterfactual endorsement, such that the Standard of Possibility is already contained within each of our practical standpoints *as a matter of fact*. It asserts this as true, and therefore concludes that all of our practical standpoints contain a normative commitment to the Standard of Possibility.

What reasons are there to think that such an empirical claim is true? The primary reason is that ordinary experience with ourselves and others might serve to suggest that it is indeed a plausible claim. Notice that this is not the claim that everyone would actually *do* what they take to be their strongest counterfactual endorsement, nor is it the claim that everyone would *admit* out loud that they are committed to this principle. It is rather the claim that each person understands and accepts the normativity of this principle, from within her own practical standpoint – that were she to reflect upon the question, "do I have most reason to do what I would most endorse," she would accept this as a normative principle for herself. For instance, ask yourself if you accept the normativity of this principle – I suspect that you do. Moreover, I believe that if most people were asked this question, the truth would be that they would accept its normativity as well.

Consider, for instance, Locke's example of a drunkard who *confesses* that abstaining from drinking further would be a greater good, yet still drinks:

> Let a drunkard see that his health decays, his estate wastes… yet the returns of uneasiness to miss his companions, the habitual thirst after his cups, at the usual time, drives him to the tavern, though he has in his view the loss of health and

> plenty, and perhaps of the joys of another life: the last of which is no considerable good, but such as he *confesses is far greater* than the tickling of his palate with a glass of wine, or the idle chat of a soaking club. It is not for want of viewing the greater good; for he sees and acknowledges it… (Locke, *Essay*: 2.21.35, emphasis added)

For my empirical claim to be true, it does not need to be the case that the drunkard actually abstains from drinking further; for people often do not actually do what they know they would endorse more. Rather, for the empirical claim to be true, it just needs to be the case that even this drunkard understands, from within his practical standpoint, that he has more reason to abstain from drinking, *given* that he *confesses* it is the greater good. And my suggestion is that *that* is a plausible empirical claim – the claim that everyone cannot help but internally recognize the normativity of this principle.

There are some things to be said for a general strategy of empirical assertion. First, consider that even if someone *claims* to deny that she has more reason to follow her stronger counterfactual endorsements, this would still not amount to a proof that the empirical claim is false – for this person could be mistaken about her own attitudes, i.e. not perfectly internally luminous, or could be engaged in self-deception, or could be telling a lie about her own attitudes. Indeed, because we cannot fully access the contents of other minds from the outside, we cannot disprove the empirical claim on the basis of observation from others: the only way to disprove it would be if my reader, herself, knew that she did not accept its normativity. Moreover, consider that even if we were to engage in a bit of foot-stomping about the truth of this empirical claim, that would still need to be assessed against the kind of foot-stomping of the realist, who must insist upon the existence of a mind-independent and brute normative fact. We might claim that, because our claim is about something more tangible, and does not face the *specific* metaphysical, epistemological, and other objections that realism incurs, we might escape with a more *benign* form of foot-stomping here.

For instance, we could offer that at least there could be some naturalistic explanation for the empirical claim, insofar as it is true.[91]

So, the first argument asserts an empirical claim as true: it suggests that the claim is plausible on the basis of ordinary experience, shows that the claim cannot be easily falsified, and furthermore notes that at a dialectical level, any kind of foot-stomping here must be contrasted with the theoretical alternative of realism.[92] But if the empirical claim is true, then it means that we all take ourselves to have most reason to pursue our strongest counterfactual endorsements – that *we* provide authoritative normativity to this principle – and this would be a non-trivial result. For the full implications of this commitment are deep: in this seemingly simple initial commitment, we would connect ourselves to a wide-ranging, objective hierarchy of mind-world potentialities. We would, through this core commitment, outsource the true contents of our reasons to a range of counterfactuals that we might have limited present information about, despite our best predictions. For note that our normative reactions to things are often involuntary, from a phenomenological standpoint: if you have never had pad thai, or never listened to Ravel, your experience of first encountering these things might surprise you, in ways that you do not currently know. That is a salient point here, for it means that our counterfactual endorsements can significantly diverge from our present predictions about them.

Let me now move onto presenting a second argument to my reader, which is meant to be independent of the first argument. This argument grounds the Standard of Possibility in the constitutive nature of normative deliberation about action, such that anytime someone enters into normative deliberation, she must be committed to the normativity of the Standard of Possibility. Because human agents do engage in normative deliberation about action, the promise

[91] I pursue some similar dialectical points in my **Empirical Argument** for morality in §5.3.
[92] Even if we cannot vindicate that *all* humans are committed to a certain kind of normative principle, even showing that a *wide swath* of them would still establish the relevant kind of objectivity *for those individuals*, and that would be an accomplishment, for objectivity is not an all-or-nothing affair (see Street 2016a: 168 & 174).

here is that the principle can be given a constitutive rooting within the practical standpoint of all such deliberators.

Here is the argument, presented first in premise-conclusion format:

1. All normative deliberation constitutively aims at realizing some future outcome that one would regard as good or worthwhile.[93]

2. All such normative deliberation about action assumes the backdrop of the most general normative question about action: "what do I have most reason to do?"

3. Because all normative deliberation about action is constitutively aimed at realizing some future outcome that one would regard as good or worthwhile, and because the guiding question of what one has *most* reason to do implies that one should prefer better outcomes to worse ones, all normative deliberation about action is constitutively aimed at realizing the future outcome that one would regard as most good or worthwhile.

4. Conclusion: All normative deliberation constitutively aims at realizing the future outcome that one would regard as most good or worthwhile, i.e. one's strongest counterfactual endorsement.

Let me now explain this argument in some greater detail.

First, I have claimed that all normative deliberation about action is aimed at realizing some future outcome that one would regard as good or worthwhile. This claim has two components to it. The first component is that all normative deliberation about action is aimed at realizing some future outcome. This, I submit, is in the constitutive structure of deliberation about *action*, because all

[93] Note that I use the term 'outcome' in a very wide sense, such that an action being realized as an end for its own sake – without any eye to its further effects – would be an 'outcome' in this broad sense of having taken place. We might also prefer to call it a 'state of affairs.' In this way, the theory is not committed to a more narrow sense of outcome, which excludes, for instance, the intrinsic value of actions being performed for their own sake, which some forms of consequentialism might exclude.

action is fundamentally future-oriented, in the sense that it is aimed at something that will take place as more time has passed. All action aims at bringing about some outcome in the future, where we should understand 'outcome' in a broad sense – including the performance of the action for its own sake. No matter what the outcome we are aiming at, because we are deliberating about what will happen next, this will inescapably be deliberation about some future outcome.

The other component of this first claim is that such deliberation about action aims at some future outcome that we would regard as good or worthwhile. This could be understood as a version of the Guise of the Good thesis, although that thesis has usually attached to claims about *desires* or *actual actions*.[94] In contrast to those claims, my claim is that all normative deliberation about action is done under the guise of what future outcome we would regard as good or worthwhile, or reason-giving. For instance, we can see this when we deliberate about ordering some menu item, or some career: we think that this is something we would endorse, for otherwise, this would have just been some whim or urge, not actual deliberation. Even if an artist is deliberating about destroying her paintings, and in some colloquial sense 'not aiming at anything good,' at the conceptual level it seems she must still be assuming some normative point to her action, insofar as it is the result of internal deliberation: perhaps it is for the sake of catharsis, or for the sake of letting out anger, or even for its *own* sake. The point is that the artist sees some point in doing it, and that is the claim here about aiming at some future outcome that one would regard as having some reason-giving status to it.

My second claim is that particular normative deliberations can only be legitimately understood against the more general, backdrop normative question of what we have *most* reason to do. For instance, suppose that you are having some particular normative dilemma, about what to order at the restaurant. This

[94] Some version of this thesis has been defended by Plato, Aristotle, Aquinas, Kant, Locke, and many others, alongside Anscombe and Raz in more recent times: see the Thomistic phrase, *nihil appetimus nisi sub ratione boni* ("we strive for nothing except under the guise of the good"). See LoLordo (2021: 24) for the point that recently, the thesis has been defended as concerning *desires* or intentional *action*; and see Velleman (1992) for a challenge to the thesis, as well as Tenenbaum (2013) for discussion.

question only makes sense against the backdrop of a general understanding that you are trying to pursue what you have most reason to do: if you receive news that your closest loved one is in the hospital, you would immediately *leave* the restaurant to go check on them, illustrating that your deliberation at the restaurant about what to order only makes sense against the background idea that you are pursuing what you have most reason to do. Similarly, your work on some project only makes sense against the backdrop of having reasons to pursue that project – and *that* again is justified in terms of what you in general have most reasons to do. More specific aims acquire their normativity from the fact that they constitutively assume more fundamental ends; and the question of which fundamental ends to adopt and prioritize just is the question of what we have *most* reason to do.

My third claim is that, when we take these first two claims together, we arrive at the conclusion that one's normative deliberation about action is constitutively aimed at realizing the future outcome that one would regard as most good or worthwhile. The first claim states that all normative deliberation about action is aimed at realizing some future outcome that we would regard as good; and the second claim is that this aim can only be understood in light of a question about what we have *most* reason to do. But because the question of what we have most reason to do implies that we should prefer internally better outcomes to worst ones – i.e. to prefer outcomes that have more reason-giving status to them to them than less reason-giving status – then the two claims, taken together, amount to the conclusion that all normative deliberation constitutively aims at future outcomes that we would ascribe most reason-giving status to. And that just is the constitutive pursuit of our strongest counterfactual endorsements, since that is equivalent to the future outcome that we would provide strongest reason-giving status to.

The conclusion of the argument is that, from the nature of normative deliberation about action, we can find that it constitutively aims at reaching our strongest counterfactual endorsements. So, return for a moment to our imagined person, who claims that she would counterfactually endorse the green curry

more than the pad thai, but also claims that she has more reason to have the pad thai. This second argument maintains that such an imagined person, insofar as she is a deliberator, is engaging in a form of self-contradiction. For when she deliberates, she constitutively aims at reaching her strongest counterfactual endorsement. When this person claims that she has more reason to have the pad thai, she just is regarding that future outcome as one that involves a stronger counterfactual endorsement than the green curry, and so she cannot also at the same time be maintaining that she would have a stronger counterfactual endorsement for the green curry, constitutively speaking. When she takes herself to have more reason to pursue X over Y, she *just is* taking herself to predict a stronger counterfactual endorsement in X over Y, all else equal.[95]

I have just offered two modes of arguing for the Standard of Possibility, each of which has an independent force. Let me now make some further observations about the Standard of Possibility, so as to avoid misunderstandings. First, let me emphasize again that we are discussing normative objectivity in the *pre-moral* domain, so if my reader is concerned about cases in which we most strongly counterfactually endorse things that negatively impact others, we will address that topic in Chapter 5. Second, I want to make clear that I am not advocating some form of *maximization* or *egoism* here. It is not maximization, because to ask the question of what one would most strongly counterfactually endorse is to ask the question of what would be one's strongest internally authoritative normative stance – but that makes no substantive assumptions about having *more* rather than *less*. For instance, we can more strongly endorse a life of having *less* money or *less* material things, and we would be under no substantive obligation to have more of anything, insofar as we endorsed having less. This is also not egoism, because endorsements have no conceptual connection to solely individualistic concern: someone's strongest counterfactual endorsement could involve spending all of one's time in the service of others (§2.3).

[95] For instance, moral considerations aside.

One might also wonder whether the Standard of Possibility forces us to think too much about the counterfactuals – and whether there is some greater wisdom or happiness in just focusing on loving the path that we have taken. The answer to such a concern is this point need not be in tension with the Standard of Possibility. The Standard of Possibility makes a claim about what, at some given point in time, we have most normative reason to do – and it explains that in terms of the strongest endorsements we could reach, from that contingent point in time. Now, it might be true that, having carved out a particular path in life, it is still the case that, twenty years ago, one had more reason to do something else. This does not undermine the point that, *having taken this path*, what we now have most reason to do is to live our lives in a way that puts ourselves behind the path we have taken: for the nature of the Standard of Possibility is principally forward-looking, telling us *now* what we have most reason to do, and that might involve most reason to be proud of the path we have taken, even though a fact still remains that we *did* have more reason to do something else.

What about babies and other animals? There are facts about what any creature would most counterfactually endorse – a baby would like being in her mother's arms more than pain, as an animal would like getting the food it wants more than starving. But the sorts of arguments I have provided about *action* and the Standard of Possibility may or may not apply to such creatures, depending on whether they capable of acting for reasons and deliberating between options. Insofar as a young infant or some animal cannot do these things, it would *still* be true that counterfactual facts about their strongest endorsements would remain. But it would be an intuitive result that we could not say that they are committed to a principle where they have a *reason* to pursue that result, as they cannot act for reasons yet, in the first place.

3.3 Endorsement and Representation

Having outlined the Standard of Possibility, let me mention two further ways in which we can make normative mistakes *by our own lights*: one that I call the **Standard of Endorsement**, and the other the **Standard of Representation**.

The first norm, which I call the Standard of Endorsement, concerns when we simply fail to realize the object of our normative endorsement. Sometimes you aim to finish some work by the end of the day, but you fail to; other times you aim to follow some new habit or practice or diet, but do not end up doing so. What explains such cases? One important reason is that making some normative endorsement does not entail, on its own, that we will end up taking some *actual action* towards its realization, or that we will have *decisive motivation* to. Consider an example:

> **Keep Working**. Kevin makes a normative endorsement to continue working on his philosophical project on this day, through his judgment that this is what matters to him most. However, he also experiences a desire to instead lie on the couch and watch Netflix. Even though he has overruled the normativity of that desire through his judgment, he still ends up spending the day on the couch and watching Netflix. As he watches, he thinks to himself with some guilt: "I am not doing what I think I have most reason to do!"

Kevin sets himself an internally authoritative normative direction through his judgment – for as I argued in the previous Chapter, *what it is* to set an authoritative normative direction for oneself is simply to overrule the normativity of some more basic desire from within the practical standpoint, say through some reflective judgment. However, the point about human psychology here is that, even though we may override the *normativity* of some more basic desire within the practical standpoint, this does not necessarily extinguish its *motivational force*, nor does it secure any *actual action*. Just because you tell yourself that you shouldn't eat the ice cream, doesn't mean you won't still be tempted to, or that you won't end up eating it. You might still eat it: and this points to the fact that there is a potential gap between *actual action* and our most

internally authoritative acts of judging, choosing, or valuing.[96] We can set, from within the practical standpoint, certain authoritative directions for ourselves, yet fail to actually *do* those things.

This dimension of human normative experience is arguably an important source of *akrasia*. On this suggestion, we can exhibit 'weakness of will' precisely because there is the possibility of a *gap* between our more reflective normative endorsements and our actions: humans who inhabit some practical standpoint are capable of telling themselves that they have most reason to do one thing, yet then not do that. As Korsgaard writes:

> We do not always do what upon reflection we would do or even what upon reflection we have already decided to do… when we do reflect we cannot but think that we ought to do what on reflection we conclude we have reason to do. (Korsgaard 1996: 104).

If it weren't for the fact that some of our attitudes had this reflective dimension to them, then this gap would either significantly narrow or not exist at all: this is perhaps what the lives of many non-human animals are like, in which their ordinary endorsements *just are* their desires, pleasures, and aversions, as the animals I am envisioning do not have the capacity to reflect upon their impulses, and to override their normative force. For such animals, there is no room for making such a normative mistake *by their own lights*, simply because their endorsements *do* end up aligning with most of their actual actions. So, the suggestion is that it is our distinctively *human* capacity for having more reflective normative attitudes that also creates this distinctively human *problem*: the possibility that we can fail to follow through on our endorsements.

[96] Locke was therefore mistaken when he maintained, in the first edition of the *Essay*, that people always *actually do* what they think is best – and this was a mistake in diagnosing the human psychological condition. Indeed, it was due to empirical objections from those like Molyneux that Locke amended his position by the time of the second edition and onwards: see again Locke's example of the Drunkard who keeps drinking, even after confessing that isn't what is best (*Essay* Book II, Ch 21, Section 35), and LoLordo (2021).

It may also be the case that we do not realize the objects of our normative endorsements, because we fail to recognize or follow through on the best, necessary, or available *means* towards realizing that endorsement. When we make an endorsement, we arguably give ourselves constitutive reasons to pursue *some* means towards realizing that endorsement, and we also give ourselves constitutive reasons to prefer *more effective* means towards realizing that endorsement, as opposed to less effective ones, all else equal. The point is suggested as constitutive of making a normative endorsement, i.e. of regarding something as reason-implying: for instance, if you claim to have reason to drink the glass of water, but deny that you should do *anything* at all towards attaining that goal – like moving your hand to pick up the water – then you will just have been failing to really maintain that you have reason to drink the glass of water. Street is right to suggest:

> Someone who "judges" that she has conclusive reason to *Y*, but who (at the same time, in full consciousness) "judges" that she has no reason whatsoever to take what she recognizes to be the necessary means to *Y*, is not making a normative judgment. (Street 2008: 227-228)

Unlike the example of drinking water, however, more complex endorsements can involve ends for which it is not obvious what the available or best means are towards realizing it – and hence there is room for failing to recognize or follow through on the means towards realizing the endorsement. The cases here can span a wide range, from needing urgently to fix a car, to trying to accomplish some vocational aim, to trying to save a person from illness but failing to follow the best medical procedures.

So, we can make normative mistakes when we fail to realize the objects of our endorsements – either because we are tempted by other competing attitudes that we have overridden from within the practical standpoint, or by

failing to pursue the available or best means to achieving the ends we have set for ourselves. This is the Standard of Endorsement.[97]

The second standard, which I call the Standard of Representation, concerns when we end up pursuing something in the real world that does not actually align with the *represented* object of our normative endorsement. To arrive at this standard, notice that many of our normative endorsements are essentially future-directed, and therefore based on some *internal representation* of the object that we have endorsed: desires, as well as many judgments and choices, are like this.[98] When we make an endorsement on the basis of some internal representation of an object, but then end up pursuing something different in the real world, we pursue something that we never gave ourselves reasons to pursue. Consider an example:

> **Desiring Gin and Tonic**. Sarah desires to drink a gin and tonic, and she ascribes a positive valence to the internally represented state of her drinking what she remembers as the taste of a gin and tonic. She sees a glass of clear liquid on the table, which looks like gin and tonic, but is in fact petrol. She then drinks the petrol and is soon at the hospital.[99]

This example, adapted from Williams, involves Sarah's giving herself normative reason to drink a gin and tonic, *i.e. gin and tonic as she has internally represented it*. That is the constitutive nature of her endorsement: since *what* she endorsed was some imagined taste and content of gin and tonic, *that* is what she has given herself reason to pursue. But then when she drinks the petrol instead, she makes a constitutive mistake: she was under a normative standard to drink gin and tonic, but she then drank something different altogether, which sent her

[97] Sometimes, we also do not reach our normative endorsements due to external factors in the world – the wind gets in our way, supply chains fail, and so on. In these cases, we might say that we do not reach our normative *ideals*. Whether or not someone makes a normative *mistake* might turn on whether the external factors were foreseeable or predictable. (On a stronger view, one could make a normative mistake in this way, even if the external factors were not foreseeable or predictable.)
[98] Unlike enjoyments or aversions that are directed towards the present experiences or sensations.
[99] Adapted from Williams (1981: 102), see also Williams (1995: 36).

running to the hospital. But she never gave herself reason to drink petrol and go to the hospital, and that is why she makes a normative mistake by the lights of the Standard of Representation.

Notice how our constitutivist explanation differs from how Williams himself explained this case, which appealed instead to how people, as a matter of fact, care about factual accuracy.[100] Williams held that normative attitudes based upon false beliefs do not ground internal reasons for action: Williams would therefore hold that Sarah never had a reason to drink the glass of petrol, since she had a false belief about its being gin and tonic.[101] But were we to ask Williams what explains the *normativity* of the requirement to have factual accuracy, he would have maintained that it is because humans have some strong *interest* in not having false beliefs: "the grounds for making this general point about fact and reasoning… are quite simple: any rational deliberative agent has in his S [subjective motivational set] a general interest in being factually and rationally correctly informed" (1995: 37).[102] But this faces two problems. First, a *general* interest in being factually informed does not entail that, in every case, the agent will maintain an interest in not having some *particular* false belief – and it therefore secures no guarantee that we can arrive at the right explanations in particular cases, insofar as the agent abandons the commitment to factual accuracy in that case. Second, the claim is also in any event empirically suspect on its own terms, because it seems that agents often *do* have interests in maintaining particular false beliefs – sometimes this even seems *advisable*, as having certain suspect in certain circumstances can help a person not harm themselves, prevent them from death, and so on.[103] Williams' justification

[100] Note that earlier, I also advanced an empirical argument for the Standard of Possibility. The point here against Williams is not that he makes an empirical assertion *at all*, but that his *particular* empirical assertion seems vulnerable to specific objection.

[101] As Williams writes, "a member of S, D, will not give A a reason for φ-ing if either the existence of D is dependent on false belief, or A's belief in the relevance of φ-ing to the satisfaction of D is false" (1981: 103).

[102] One could alternatively try to interpret Williams as saying that all belief inherently seeks truth – that it is in the nature of belief to do that – though it doesn't seem to me that William suggests that in the text himself.

[103] In moments of great despair, we sometimes need to believe certain things just to motivationally sustain ourselves, even if they are not true. Williams seems to concede as much:

therefore seems too weak to support his general claim that attitudes based on false beliefs do not generate reasons, especially in cases where the individual claims not to care about factual accuracy.[104]

Our explanation promises to be more precise and secure: we explain Sarah's not having a reason to drink petrol, not through some general interest she has in factual accuracy, and not through a brute normative requirement, but from the fact that she never constitutively gave herself reasons to drink petrol, given that she endorsed an internal representation of gin and tonic, which tastes and has effects that are nothing like petrol.[105]

Consider now an example of the Standard of Representation in practice:

> **Aspiring Celebrity**. Ashley has chosen to aspire towards becoming a famous pop artist, on the basis of some internal representation of what that life is like: imagining that, along with getting constant attention from others, she would also feel fulfilled and flourishing. Upon becoming a famous pop artist, however, Ashley realizes that it is lonely, isolating, and crushing in its pressure.

The reality of such a case is bolstered not only by common observation, but by psychological studies of actual celebrities that report this exact pattern of initial desire, coupled with actual disappointment.[106] The Standard of Representation

---

"There could be a case of somebody who had an overwhelming need to be deceived; and if his relations to reality were so poorly negotiated that he actually needed to believe what was false, then perhaps he would have reason to acquire false beliefs – in that particular respect" (1995: 37).

[104] An alternative possibility is to read Williams as a realist, in our sense of the term: that he sees the normativity of not having false beliefs as something like a *brute* requirement. Perhaps Williams himself thought that a 'procedural' requirement of having accurate beliefs was somehow more meta-ethically benign than a directly substantive requirement, like Markovits (2014) maintains. But as I suggested in §1.4, even one claim about brute normativity is what moves a theorist into realist territory, and so Williams then would face all the same objections to realism that we have encountered (§1.2).

[105] A requirement along these lines is mentioned in Sobel (2009: 344) and Dorsey (2017b: 207-208). However, while Sobel and Dorsey frame their points in terms of providing some theoretical *rationale* for idealization, they do not emphasize the same kinds of constitutivist points that we do: for instance, Sobel merely wants to claim that his proposed "rationale is *compatible* with the ambitions of subjectivism" (2009: 339, emphasis added), but not that it is *justified* by mind-given or constitutivist points. Moreover, these idealizing requirements also involve having a 'perfect forecast' (2009: 345), which is arguably stronger than what we are after here.

[106] For instance, according to a recent psychological study of fifteen celebrities, the authors concluded: "From an initial desire to become successful, the celebrity experiences personal

offers a diagnosis of these cases: when we work hard at achieving something on the basis of some internal representation, we open the risk that we could end up reaching something that never resembled the picture that we had in our heads. For what Ashley endorsed, and therefore gave herself reason to pursue, was a life that involved feelings of fulfillment and flourishing; but instead, what she got was a life that involved loneliness and isolation, which she never bestowed positive normativity upon.

Let me close this section by contrasting the Standard of Representation with the Standard of Endorsement. I mentioned that both standards are, in some broad sense, about failing to meet our normative endorsements. But these two principles hone in on different reasons for such failures. The Standard of Endorsement takes effect when we *accurately pursue* the object of our endorsement, but simply fail to realize it, for *akratic* or other reasons. By contrast, the Standard of Representation takes effect *not* when we fail to succeed in attaining some object, but when the object we attained was not *accurate to* what we had represented in our minds when we made the endorsement. So, to have a metaphor: both standards concern failing to hit our target; but the Standard of Endorsement is like *aiming correctly* but *shooting poorly*, while the Standard of Representation is like *aiming poorly* but *shooting correctly.* Taken together, the normative imperative is that with our endorsements, we must both aim at them accurately, and execute successfully.

3.4 Substantive Upshots

How do these principles now help with the substantive challenges that we outlined at the outset of this Chapter? I have already provided some examples of how the Standards of Endorsement and Representation set norms against which we can make normative mistakes, such as when Kevin fails to work on

---

confusion and a loss of ownership of life in a depersonalizing 'entitization' process, in which participants reported feeling like a thing rather than a person of unique character. Immediately upon entering the sphere of fame, relationships with 'self' and 'other' are profoundly affected" (Rockwell & Giles 2009: 185).

his project, or when Ashley pursues fame that does not provide her what she imagined. Let me now discuss how the Standard of Possibility – the deepest of the principles we have outlined – applies to substantive cases.

We can begin with the grass counter. When the grass counter claims to find great pleasure from counting grass, the normative question is no longer exhausted by the fact of his actual pleasure: for the Standard of Possibility tells us to ask whether the grass counter *would* endorse doing something else more, if only he tried out that alternative, potential activity. For instance, in Rawls' imagination, the grass counter has mathematical abilities – so we would need to ask whether the grass counter might have a stronger counterfactual endorsement in a vocation of discovering profound mathematical truths (and not just solving problems for a fee). Insofar as the grass counter would find greater holistic endorsement in that alternative – or indeed, insofar as the grass counter would find greater holistic in *any* other alternative – then he will have stronger reason to pursue the alternative, according to the Standard of Possibility. For, by his own lights, he will have committed himself to the principle that he has more reason to pursue his stronger *counterfactual* endorsements. The lesson with the grass counter is one that pertains to all of our lives: though we might think that what we are doing in the present is decent, or acceptable, that does not mean we are doing what we have *most* reason to do. For the normative injunction that we face, according to the Standard of Possibility, is to discover our own highest normative potentialities: to be doing those things that we would *most* find worthwhile or meaningful or important, once becoming acquainted with those activities. It is therefore a principle that gives content to the idea of living our best lives.

But here is the flipside that we also need to insist upon: if we imagined a person for whom there was *no* stronger counterfactual endorsement besides counting grass, the theorist should at this point relent and accept that this person *does* have most reason to count the grass, at least prior to moral

considerations.[107] For we really need to think about who we have imagined: we have imagined someone who, from his internal perspective, sees so much value in the act of counting grass, such that there is no other activity which could give him a stronger sense of internal enjoyment, positive judgment, or other form of authoritative endorsement. Indeed, we might have no idea *what* this person sees in grass counting, but almost as a definitional point about the assumption that she has her strongest counterfactual endorsement here, *she* must see something it: perhaps it is the beauty of seeing blades go by one by one, or the ascetic value of focusing on one activity, or the like, that explains his love. The point is that, at *this* point, the theorist should take on a stance of openness and interest in his most unique point of view on the world, rather than to continue to cast normative judgment from the outside.[108] After all, many of the activities that we humans actually pursue, in less far-fetched examples, involve commitments to things that a wide range of other humans do not understand: these are the tea tasters, Everest-climbers, tightrope-walkers, classical-music lovers, and for that matter, *philosophers* of this world. For many people do not get why philosophers are so invested in philosophy, but that does not count against our having reasons to pursue philosophy, given our strong normative relationship with it, from our practical standpoint. I therefore offer the substantive objector a dilemma: either our theory *can* vindicate that the grass counter *does* have more reason to do something else, if she has a stronger counterfactual endorsement elsewhere; or otherwise, we should relent on the intuition that demands we say something more about him in the first place. Either way, the point is that the objection can be diffused.

The Standard of Possibility has deep implications for our reasons. A young girl in an ancient or oppressive society might not realize that reading

[107] After the arguments for moral objectivity of Chapter 5, we may find that the grass counter – in leading a life that is insensitive to what more could be done to help others – is not living up to what she has most reason to do.

[108] Compare Street (2009a) for points in a similar spirit. I must emphasize here that we are not only trying to imagine some being who loves grass counting, but who loves grass counting in a counterfactually tested sense – there is *nothing more* she could love more. That would indeed be a rare type of being.

books is a worthwhile pursuit for her, as she has never been given an opportunity to try: but this does not mean that she has no reason to read, as she most certainly can find some strong counterfactual endorsement in the joys of learning. Someone who is about to walk down one vocational path, under pressure from societal expectations, might not realize that there is some other choice that could better serve her talents and passions in a holistic sense – and this person would be going against her stronger reasons, in light of the Standard of Possibility. The Standard of Possibility renders normative for us an expansive range of objective mind-world potentialities, and it tells us to find the deepest ways in which we could come to relate to the world, given our contingent point in time and space.

Conclusion

This Chapter has argued that a mind-given theory can vindicate that we can make normative mistakes in a number of ways, and hence that we are beholden to certain objective pre-moral norms. The Standard of Possibility connects us to an objective hierarchy of mind-world potentialities, telling us to pursue our strongest *counterfactual* endorsements; the Standard of Endorsement tells us to live up to our own highest endorsements, which can often be hard to do in practice; and the Standard of Representation tells us to ensure that our actual pursuits are aligned with the true, internal nature of our endorsements. While I have not claimed that these principles exhaust normativity objectivity, I have claimed that these principles can go far in explaining our ordinary intuitions about normativity, and that this allows the mind-given theorist to achieve a layered and robust picture of what is worth pursuing in human life.

In the next Chapter, I turn to discuss some deeper nihilistic concerns towards a mind-given theory.

4 Nihilistic Concerns

4.1 Introduction

We have been developing a theory that began with the idea that normativity arises from the mind – and as we have seen, there are a number of important questions to be explored for such a position. Now, we turn our attention to a deeper set of concerns, as to whether a mind-given theory is *nihilistic*.

Nihilism is a broad term, so let me distinguish at least three types of nihilistic concerns. First, many of us have a deep yearning for value to be *timeless*. We would like to think that sunsets will remain beautiful, even if human minds disappear from the world; but the mind-given picture seems to imply that the beauty of the sunset will fade once there is no longer anyone in the world to appreciate it. Connected to this, and more broadly, the mind-given picture seems to threaten the ordinary idea that some things have *value*, even if they are not being *actively valued*.[109] Second, some have claimed that a mind-given picture entails that there are only *descriptive* facts about human psychology, rather than genuine *normativity*. Parfit writes, for instance, that "normativity is either an illusion, or involves irreducibly normative facts" (2011b: 267). Third, some have objected that mind-given normativity is not *real*, in some important sense – that there is something objectionable to the thought that normativity is a *mere* phenomenon of the mind. Iris Murdoch, for instance, in advancing a Platonism about the Good, writes that "it was better to know what was real than to be in a state of fantasy or illusion" (1970: 62).

Nihilistic concerns along these lines have been acknowledged by both critics and champions of the mind-given picture. Parfit writes that a mind-given picture is "bleak" and entails that "nothing matters" (2011a: 107), insisting that such a view would mean that he wasted much of his own life (2011b: 367). Similarly, Street writes that, on an antirealist view like her own, "the threat of

[109] In this Chapter, I use the notion of 'valuing' in a broad sense, as referring to any possible act of normative endorsement, from simple enjoyments to complex judgments.

nihilism never goes away" (2017: 123), and Mackie concedes that the rejection of mind-independent values can lead to "a feeling that nothing matters at all" (1977: 34).

I believe that a mind-given position need not entail these nihilistic conclusions. In this Chapter, I consider each of these concerns, and I develop responses to them. I begin with the deepest concern: that a mind-given picture cannot capture that value has a timeless quality to it, with an independence from what is actually being valued. In response, I develop a view in which we can understand value in terms of normative *potentialities*, and I argue that this allows us to find a set of *objective* and *unchanging* facts that can mirror the realist's mind-independent values (§4.2). I then address the objection that there are only *descriptive* facts on the mind-given picture, arguing that this objection misses the fact that there is genuine *normativity* from within the practical standpoint (§4.3). Finally, I consider the claim that mind-given normativity is not *real*, and I argue that this claim rests upon a mistaken understanding of the place of conscious experience in reality (§4.4). I conclude that a mind-given theory need not be nihilistic: to the contrary, it can recover a rich and exciting picture of the place of normativity and its potentiality in the world.

4.2 Value and the Import of Potentialities

Let me begin with the deepest problem, which is that the mind-given picture seems to provide no anchor for there being *value*, apart from things that are *actually being valued.* There are two dimensions to this problem. First, this seems to threaten our classical conception of value: we typically think that many things have value, even if they are not being actively valued. For instance, we might think that the Mona Lisa and Beethoven's Fifth Symphony would retain a certain kind of value, even if they suddenly faded from our collective memory or active appreciation; similarly, we might think that objects or artifacts in distant lands have value, even if we have never had the chance to see and appreciate them for ourselves. On the classical conception of value, when

something *has value* or *is valuable*, there then seem to be normative reasons to value it.[110] The problem is that the mind-given picture seems to make no space for such value, since all it provides us are facts about what people *actually value*. Outside of that, the picture seems empty, as it sets aside the existence of value that is ontologically independent of the mind.

The problem at hand can be witnessed in a famous exchange between Mill and G. E. Moore. There is a passage in Mill, in which he can be read as conflating 'desirability' and 'actually desired':

> The sole proof capable of being given that a thing is visible, is that people actually see it. The only proof that a sound is audible, is that people hear it; and so of the other sources of our experience. In like manner, the sole evidence it is possible to produce that anything is *desirable*, is that people do *actually desire it*. (Mill 1863 [2015]: 148, emphasis added)

This passage drew the philosophical ire of Moore, who pointed out that 'desirable' means something different from 'actually desired':

> The fact is that 'desirable' does not mean 'able to be desired' as 'visible' means 'able to be seen.' The desirable means simply what *ought* to be desired or *deserves* to be desired… 'Desirable' does indeed mean 'what it is good to desire'; but when this is understood, it is no longer plausible to say that our only test of *that*, is what is actually desired. (Moore 1903: 66-67)

Moore added that "the fallacy in this step is so obvious, that it is quite wonderful how Mill failed to see it" (1903: 66).[111] So, we have the question at hand: what

[110] This should be understood as the proposition that there are *pro tanto* normative reasons to value it: that if some work of art, landscape, object, or activity has value, then its value *counts in favor* of our coming to appreciate that value, at least to some extent. Valuable might also be taken to mean something like *worth* valuing, or that we *should* value it (see Crisp 1997: 75 and Moore 1903: 66, who respectively make these points in terms of desirability).

[111] But note that Mill may only have been seeking to make an epistemological point, rather than a grounding or conceptual point. See Crisp (1997: 72-77) for helpful discussion.

can a mind-given picture say to the charge that it leaves no space for value (or 'desirability') outside of valuing (or 'actual desiring')?

I said that there were two dimensions to this problem, and the second one is even deeper: it has to do with the fact that we would like to vindicate that value has a *timeless* or *everlasting* nature to it. For there is something profound that seems lost in the conclusion that value will recede from the world, alongside the disappearance of sentient standpoints. Consider an example:

> **Sunset.** You and your friend are walking on a boardwalk, and there is a sunset that puts both of you in awe. Your friend turns and asks you: "Do you think that the sunset would still be beautiful, even if humans went extinct one day, and there was nobody here to appreciate the sunset?" Having recently become skeptical of the existence of mind-independent values, you feel compelled to answer *no*: "the value of the sunset will disappear the moment its valuers disappear." With that, you both feel some sadness inside.

Such an exchange will be familiar to many of us, and it could be switched to match a variety of contexts or personal experiences. For when we reflect upon our deepest encounters with what we find beautiful or meaningful – moments in which we sense that there is some great value in the world – we tend to think, or at least would *like* to think, that there is something metaphysically profound that is happening: that we are getting in touch with something that is larger than ourselves.[112] It was Plato who memorably captured this impulse in his conception of the Forms: "eternal, nonphysical, quintessentially unitary entities… standing immutably in the nature of things as standards."[113] There is something quite soaring to the thought that, when we encounter the beauty of the sunset, or the sublimity of some music, we are encountering the Form of

[112] Many of my readers will also think about God here. But we must recall a dialectical point, which is that this is a project that explores what can be said at a theoretical level *without* assuming that theism (or realism) is true (§1.2). Obviously, however, much could be said here about God and theism, which will need to wait for another day.

[113] See Cooper (1997: xiii). The Forms can be found discussed in a variety of Plato's works, including the *Phaedo*, *Republic*, *Phaedrus*, and *Symposium*.

Beauty itself – a timeless standard that will long outlast ourselves. It can be comparably dismal, by contrast, to think that beauty will fade from the world, once there are no more creatures around to see things as beautiful.

The responses that mind-given theorists have offered to this situation have not approximated the sense of depth and wonder that Plato's conception had. Nietzsche, for instance, imagined that after the cooling and congealing of our sun – and hence the loss of sentient life as we know it – "nothing will have happened."[114] Mackie also gave acknowledgment to the sense of existential loss that could accompany the rejection of realism:

> The denial of objective values can carry with it an extreme emotional reaction, a feeling that nothing matters at all, that life has lost its purpose… the abandonment of a belief in objective values can cause, at least temporarily, a decay of subjective concern and sense of purpose. (Mackie 1977: 34)

Mackie's solution to this predicament was essentially just to *soldier on*: he reminds us that "the lack of objective values is not a *good reason* for abandoning subjective concern for ceasing to want anything," and notes that such a feeling of loss "does not *follow*" (1977: 34, emphasis added), as if existential sentiments were amenable to straightforward logical correction. Street acknowledges the same basic problem:

> The whole idea behind antirealism is that things matter, but only ultimately in virtue of your thinking they do. The *causal* threat of ceasing to think things matter never goes away, and there's a kind of vertigo involved in recognizing that value is there only as long as you think it is and that you could slip into a state in which you stop thinking there is. (Street 2017: 148)

As a solution, Street offers us some practical advice that goes a bit farther than Mackie: "it seems to me the only measures we can take against *that* possibility

[114] See Nietzsche, *On Truth and Lies in a Nonmoral Sense* (1873).

are causal – getting enough sleep, eating right, cultivating a wide and deep range of interests, and so on. We should fear the view that nothing matters… as a state of mind we might fall into – false as long as we don't fall into it, but true as soon as we do" (2017: 148). Street's practical exhortations may be helpful, but the question remains as to whether the mind-given theorist can offer a deeper theoretical solution.[115]

Let me now move to explain the solution that I have in mind. On the suggestion I want to advance, we can understand the *value* of some object in terms of the *potentialities* of how that object *could* be valued by some possible mind.[116] This proposal would promise to accomplish two things: (i) to provide an anchor for value in potentialities that would allow us to speak of value as independent of what is actually being valued, and in a way that aligns with our pre-theoretic intuitions; and (ii) to point our attention to a set of facts – facts about potentialities – that are themselves *objective*, *unchanging*, and *everlasting*, much like Plato's Forms.

On the mind-given picture, the sunset does not have some mind-independent, ontological property of value; it is something that is seen as beautiful from the perspective of a human mind. But the rejection of some mind-independent property here does not mean that there is nothing left to be said about the sunset: for there still remains a certain kind of normative *potentiality* that surrounds it. The sunset has the *potential* of being seen as beautiful, in virtue of its descriptive characteristics, coupled with fact that there are possible minds – minds like ours – that could come to see the sunset as beautiful. That is a *fact* about the sunset that will never change, because the nomological possibilities in our world will never change – the possibilities of mind-world relations that inhere in virtue of the basic metaphysical structure, constants, and laws of our universe.[117] Our universe is one that *permits* the possibility of conscious minds

[115] For a sense of the nihilism at hand, see also Williams' discussion of Nietzsche and Bertrand Russell in "The Human Prejudice" (2006: 137-138).
[116] I use 'object' here in a wide sense, to include things like activities, landscapes, works of art, and so on.
[117] This appeal to nomological possibilities is be contrasted with the stronger notions of metaphysical and logical possibility, which concern possibilities in which the basic laws of

like our own, as well as the possibility of objects like the sun – and given these basic conditions, the *normative potentiality* that inheres at the intersection of some possible mind and some possible object will always remain the same. Indeed, the space of possibility is filled with a hierarchy of mind-world relations, all of which involve potential standpoints from which something could be seen as deeply beautiful, meaningful, or valuable. Such normative potentialities are a powerful tool for the mind-given theorist, because every such potentiality does not ever lose an essential reference to some possible *mind*, alongside some possible *object* – and in this way, a potentiality is different from the existence of an ontological property of value that is entirely mind-independent.

To speak more precisely, every normative potentiality has some *indexical relation* to some particular mind, or some set of minds. For instance, the sunset is something that could be seen as beautiful by *the human mind*, though it is not something that other animals like pigs are (to my knowledge) likely to see as beautiful. Similarly, *rolling in the mud* is an activity that could be enjoyed by *the mind of a pig*, but not so much by the human mind. Here, we are providing examples in which we can make generalized claims about the mind of some species, based on shared counterfactual dispositions between the different individuals of that species. But we might be more specific in other cases: given that each of our particular minds might show variation in other examples, there can be even more limited and precise claims about the normative potentialities that surround particular minds *within* a species. For instance, Sarah might be able to come to see a whole lot of value in certain jarring or esoteric forms of art, while Cayla might not share those same counterfactual dispositions; and then there would be different claims about the normative potentialities that surround those distinctive minds.

Now, my proposal is that we can understand the *value* of some object in terms of its potentialities of being valued – and I want to suggest that this allows us to meet both dimensions of the nihilistic problem. First, this seems to offer a

---

nature could have been different. If one wished to take a stronger view of possibility, the same essential point could be maintained here, though in an amended manner.

conception of value that aligns with much of what we were looking for. When we say that Beethoven's Fifth Symphony has value for humans, what we mean is that it is a piece that humans could come to experience a great deal of beauty in. What makes Beethoven's Fifth Symphony much *more* valuable than Kevin's Toy Symphony is that there are much stronger potentialities of valuing that surround it: as much as humans may try to enjoy Kevin's Toy Symphony, it just is not going to create a comparable experience of beauty as Beethoven's Fifth Symphony. This point about the value of Beethoven's Fifth Symphony, as being grounded in its potentialities of being valued, is independent of how many people are *actually* valuing it, as a matter of fact. Even if everyone forgot about the piece, or disappeared from the world, these potentialities surrounding the Fifth Symphony would remain the same, because there would still be the same *possible minds* that could see beauty in it. This is what allows us to achieve an important distance between value and actual valuing, as we found in the classical conception of value.

But moreover, let me suggest that we can vindicate the idea that, in virtue of the fact that the Fifth Symphony has value in this sense, we thereby have normative reasons to come to appreciate it – another feature of the classical conception of value. Why is this the case? It is because I argued in the previous Chapter that we have normative reasons to pursue our strongest counterfactual endorsements – the endorsements that we *could* make from different potential standpoints. In §3.2, I argued that this was a conclusion entailed by existing commitments from within our practical standpoints. But notice that when we speak about the value of some object, as *indexed* to some mind, we simply are speaking about certain counterfactual endorsements that are achievable for that mind. Insofar as these are counterfactual standpoints that would involve a strong endorsement of an object, then we give ourselves normative reason to realize those standpoints and to come to appreciate what they reveal to us about the object. So, for instance, suppose that I have never listened to Mahler's Eighth Symphony, but that I would come to appreciate it as beautiful, if I simply put it on the stereo. On our conception of value, the Eighth Symphony would thereby

have a value, as indexed to my mind, since there is a potentiality of valuing for me – and at the same time, I would have a *reason* to listen to it, in virtue of this counterfactual endorsement. So, when we couple this conception of value with our previous argument about counterfactual endorsement, we can arrive at the elegant result that value (as potentiality) entails reasons to value, since we have reasons to pursue our counterfactual endorsements.

This conception of value also seems to align with our pre-theoretic intuitions, if we consider some examples. We tend to think that natural landscapes are valuable – and, on the present proposal, this is because they are the types of things that human beings could come to find beautiful, upon seeing them. Similarly, flowers might have value for human beings because of the potentialities of our coming to see them as beautiful; but they might have a different kind of value for bees and insects, who could come to value them as sources of food or pollination. Moreover, this conception explains why some activities are only valuable *for humans who are disposed to like that activity*, such as for intense forms of mountain-climbing or tightrope-walking; why certain other activities are valuable *for particular animals* but are not also valuable *for human beings*; and why, generally speaking, there is much that is valuable which we are currently failing to value. In this way, this conception of what is valuable also brings us closer to its etymological roots: as what is value-*able*.[118]

To summarize, we can vindicate a number of elements in the classical conception of value through an account that grounds value in potentialities of valuing. First, this provides some anchor for value outside what is actively being valued; second, we can reach the conclusion that such value entails normative reasons to come to appreciate the value; and third, it makes sense of our pre-theoretic intuitions. This is not to say that every aspect of a classical conception of value will be carried over – for instance, perhaps some people will not be used

[118] One might also compare a suggestion from Sidgwick, when he writes that 'desirable' means something like "what *would* be desired" under certain conditions, including with an awareness of the object (1907: 110-111).

to understanding value in terms of what is indexed to particular species or individuals. But we can be prepared to suggest that we should not lament these shifts, for there would in any event be something dogmatic in thinking that claims about value could be made without any reference to what kind of creature could appreciate that value, or in a way that privileges only the human mind – for many of the other animals also experience the world in normative terms. So, depending on one's starting presuppositions, our proposal can be understood as either *elucidating* what lies behind an already familiar notion of value, or as *reforming* our ordinary conception of value into something that aligns with the other philosophical commitments we would like to maintain.[119]

What about the deeper dimension of the problem, which concerned our longing for value to be something timeless and everlasting? Well, think for a moment about what Plato's Forms were: these were invisible, ethereal, abstract, and unchanging facts that set standards for normativity. What I would like to suggest to the reader is that normative potentialities are not so different: they are also invisible, ethereal, abstract, and unchanging facts that set standards for normativity.[120] Our universe is full of normative potentialities, including many that we are not aware about – for a normative potentiality inheres every time there is some possible mind-world relation in which normativity could be experienced. Because normative potentialities inhere in the basic nomological features of our universe, they are timeless and everlasting – the same normative potentialities were there at the first moment that our universe took on the fundamental metaphysical structure it has, and these potentialities will continue to be the same, long after valuing creatures are gone. Appreciation of this fact can lead to a certain kind of reverence or awe for the fact that conscious life – and moments in which we relate to the sunset in a meaningful fashion – are

[119] For instance, Lewis proposes an "analytic" and "naturalistic" conception of value that might be understood as falling under either of these aims (1989: 113), and Street speaks about developing a "reforming" conception of value (2011: 31).

[120] One difference, of course, is how normative reasons for creatures arise in light of these facts: on our view, normative reasons arise when there are creatures around who then *constitutively* connect themselves to these normative potentialities, thereby rendering them into reason-giving ones.

possible in our universe at all. So it seems to me that reflection about potentialities can provide some existential solace to those who were moved by these sorts of concerns in the first place: we can capture the sense in which Plato's vision was right, of there being eternal and unchanging facts about what humans can and ought to value, simply by grounding that sense in a realm of potentialities, rather than a realm of ontological entities.

Indeed, because normative potentialities extend beyond just *actual* objects, and also obtain in virtue of there being *possible* objects, this view can take us even further than traditional realists like Moore and W. D. Ross could. For on their view, so long as objects *and their properties of value* continued to exist, then there would be value in the world; but because objects *and the properties of value that attach to them* can disappear, this ultimately results in a view that does not make value deeply permanent. For instance, Ross recognized that if the primrose *and its attribute of fairness* did not exist, then there would be no fairness in the world:

> Now it is true that the primrose could not be fair or yellow if nothing but it existed. *It could not be fair if its fairness did not exist*, nor yellow if its yellowness did not. (Ross 1930: 77, emphasis added)

But what happens once the sun loses its heat one day? Or what happens when the sun disappears altogether, along with the property of beauty that supervened upon its natural characteristics? Even an appeal to *mind-independent* properties of value cannot establish that these properties will never disappear, for insofar as the properties supervene on fragile natural facts, these mind-independent properties can disappear as well. This is what makes a view about *potentialities* even more secure than this kind of traditional realism: for potentialities are grounded in the possibility of some mind, as well as just the *possibility* of some object. And so, even after the sun loses its fire, there will still be a normative potentiality, in virtue of the fact that a sun was possible, and a human mind,

seeing it from a certain distance and angle, was possible – and that fact about this everlasting normative potentiality is one worth marveling at.[121]

I have suggested that a view in which value is understood in terms of normative potentialities can help assuage nihilistic concerns associated with the idea that actual valuing might disappear. Let me now apply some of these points through a reflection upon musical beauty. I want to suggest that our account of value, as grounded in potentialities, can offer a compelling explanation for why the beauty of a work of music or art is not always appreciated upon first debut. Consider Tchaikovsky's First Violin Concerto, which is now beloved by many and performed regularly around the world – a composition that certainly seems to have a plausible claim to being valuable. When the First Violin Concerto premiered, however, critics did not react to it favorably: the well-known Austrian critic, Eduard Hanslick, described it as "long and pretentious," saying that it "brought us face to face with the revolting thought that music can exist which stinks to the ear"; moreover, the concerto gained a reputation as unplayable.[122] But what explains the fact that the First Violin Concerto is now considered a beautiful masterpiece, when so many people did not value it at first?

If value has to do with discerning mind-independent normative properties through, say, a faculty of intuition, then we must claim that critics of the time suffered some sort of collective, meta-ethical malfunction at the premiere: although the Concerto had some timeless property of value all along,

---

[121] It seems to me that some more recent realists might also face a challenge in capturing the *timelessness* of value, in light of their remarks suggesting that a world without creatures may also not contain value. For instance, Scanlon can be read as maintaining that value is to be explained in terms of reasons, where a reason only exists when there is an *agent* there to have the reason – after all, "is a reason for" is a four-place relation R (p, x, c, a) in which the agent herself is one of the necessary constituents of this relation (2014: 31), and Scanlon emphasizes that "the existence of such agents is a presupposition of the practical domain" (2014: 22). Scanlon's view might therefore be read as implying that, once there are no more agents, there will also be no more value. Similarly, when Alison Hills outlines a Kantian value *realism* in which "rational nature is valuable whether or not we recognize that it is and choose to treat it as valuable" (2008: 188), she adds that this is *not* the position that "values would exist in a world devoid of creatures who see and respond to them" (2008: 188-189). These views – insofar as they are also concerned about timelessness – may face questions similar to the ones that we have been addressing here.

[122] See Dotsey (2018) for the quote, and Tchaikovsky ("Letter 1916 to Nadezhda von Meck") for the latter point.

this was not intuited by a large number of people. This explanation would require effort, in that we would need some compelling story of this collective failure of epistemological discernment or intuition. Instead, on the account being advanced here, we can say that critics of the time were *conditioned* into seeing music in certain ways – they occupied a certain kind of *normative lens* (§2.5) – and were therefore out of touch with different possible ways of appreciating the concerto. If they were only willing to adjust their normative lens – to look at the concerto from a new set of presuppositions, or with a new ear – then they could have come closer to a higher set of counterfactual endorsements. The higher potentialities of hearing beauty in the Concerto existed even at the time of its debut, in virtue of the mind-world relations that were possible between the human mind and the Concerto; they simply were not realized at the time. This explanation of why the Concerto was not appreciated at first, but is now, serves to account for both (i) why the Concerto has a stable value over time, yet (ii) how there can be a collective failure to actually appreciate what is of value. On this view, it was not that there was some collective failure of intuition towards an unchanging mind-independent property; rather, there was a collective failure in seeking out a certain potentiality of appreciation.

There are many things in the world that are valuable, but which we may have not realized are valuable. For we might not know all the potentialities of seeing beauty or goodness or value in things. So, it behooves us to be attuned to the potentialities of valuing, and hence to value.

4.3 Mind-Given Irreducibility

We turn next to the concern that, on a mind-given picture, there are only *descriptive* facts about human psychology, rather than genuine *normativity*. Parfit makes this objection, when he writes:

> When Subjectivists claim that some things matter to particular people, they mean only that these people *do* care about these things. That is not a *normative* but a *merely psychological* claim. (Parfit 2011a: 107, emphasis added)

Parfit adds that "we hoped that philosophers, or other wise people, would tell us more than that" (2011a: 107). He also makes this objection in his discussion of Williams, writing that "Williams did not understand the distinctive concept of a non-psychological *purely normative* reason" (2011b: 435). Parfit continues:

> Williams thinks of reasons, not as facts that *count in favour* of our having some desire or acting in some way, but as facts that might *motivate* us. (Parfit 2011b: 435)

> I assumed for many years that Williams had some purely normative beliefs. I failed to see that Williams's claims about reasons, and about what we ought to do, are really psychological claims about how we might be motivated to act. (Parfit 2011b: 452)

Parfit's claim is that, on a mind-given position like Williams', there are only descriptive facts about states of psychological motivation, rather than facts about *normativity*. Parfit suggests that, for there to be *normative* facts, there would need to be facts about what *counts in favor* of some kind of action or desire, not just observations about what people might be motivated to do. This can also be stated as the objection that there is no *action-guidingness*, on a mind-given theory.[123]

In order to respond to this objection, we need to ask what it would take for there to be *normativity*. If Parfit simply equates normativity with the idea that there are mind-independent standards of normativity, then by definition, a

[123] It is an interesting question whether Parfit interpreted Williams correctly or not. For instance, Williams insists that his view about reasons "has normative force," because it can allow us to "improve or correct" an agent's "beliefs of fact and his reasonings" (1995: 36). But there are outstanding questions here, including if an agent has no necessary reason to care about factual accuracy (see §3.3). See Street (2017) for discussion of Parfit's interpretation of Williams.

mind-given picture will not capture normativity in *that* sense. But that does not seem to be the ordinary meaning of normativity: as Parfit points out, normativity involves the more general idea that some things *count in favor* of doing certain things. Broadly speaking, normativity seems to involve there being something distinctively *non-neutral* to our encounters with the world; as I said in the opening Chapter, normativity involves there being some *directionality* to our lives, and this is what makes normativity different from empirical or scientific description, which is neutral.[124] Telling someone that the mountain is a certain height is to make such a neutral observation about it; but to tell someone that it is *valuable* or *beautiful* is to enter into normative territory.

When we are interested in normativity in this general and familiar sense, we can respond to Parfit that there is plenty of normativity from *within* the practical standpoint – for life from the driver's seat of being human involves far more than mere *description* of the world. When we experience the symphony *as good* from the standpoint of conscious being, or the pain *as bad* from our standpoint, we are having *normative* experiences – we are taking certain features of the symphony to *count in favor* of our continued listening of it, or certain features of the pain to *count in favor* of our escaping it. Although we can certainly make descriptive statements about human psychology *from the outside*, things *become normative* once we take on the actual standpoint of a human person. To illustrate this point, consider the following example:

> **Sophie at the Doctor.** Sophie is a child at the doctor's office, and she is undergoing some excruciating pain. Sophie tells the doctor: "I don't want to be in this pain, it feels so *bad*!" The doctor then writes down on his notepad: "Sophie is in a psychological state in which she is motivated to not be in the pain."

As Parfit points out, it is true that the doctor can make a descriptive statement about Sophie's pain on his notepad, from an observer's standpoint. But what

[124] We also began this dissertation with the idea that normativity is a *basic* concept, but which can be given some elucidation through analogy or explication (§1.1).

Parfit's diagnosis leaves out is that there is a first-personal dimension to Sophie's experience that is *irreducibly normative*. Sophie's experience is irreducibly normative, in two important senses: first, the normative *concepts* that Sophie uses to represent her experience are basic and irreducible *as concepts* ("the pain is *bad*!"); and second, the very nature of her experience has an irreducibly normative character to it, in that the nature of the experience cannot be reduced to merely *neutral* description.

Both of these points are important. First, normative concepts can be irreducibly normative, even if they are understood as things that we use from within the practical standpoint. When we say that something is *good* or *bad*, we cannot say what we mean in a more basic way except through those particular normative concepts – concepts which cannot be reduced in their meaning to other concepts, especially not descriptive concepts. For this reason, both realists and non-realists have agreed that normative concepts are irreducible: Parfit writes, for instance, that a normative reason is "indefinable" (2011a: 31), and Scanlon writes that it is "primitive" (1998: 19). Street writes as well that "normative concepts are irreducible and indefinable" (2017: 126). And note that the proposition that normative concepts are irreducible, as concepts, is distinct from the proposition that there are brute and mind-independent *facts* about normativity.

But second, our experiences of normativity themselves also have an irreducibly normative character, in the sense that there is *something it is like* to experience the world in distinctively non-neutral terms, which cannot be reduced to merely neutral description. When you experience the pain as bad, you cannot capture the nature of that experience – which is a phenomenological *fact* – through the kind of neutral description that characterizes scientific observation from the outside.[125] Indeed, as an even deeper point, the core nature of our

[125] Again, to say that normative experience has an *irreducibly normative* character is *not* to say that there are irreducibly normative *facts*, i.e. normative facts which are mind-independent and brute; it is simply to say that the *character of our experience* cannot be reduced to merely neutral description. Our points here can all be understood as fundamentally reliant upon the existence of minds; and moreover, we are *not* claiming that the very existence of minds is not open to further explanation.

normative experiences are often *ineffable*, in that even our use of *normative* concepts does not exhaustively capture what it is like to undergo those normative experiences. When I tell you that the symphony is *beautiful*, I am barely scratching the surface of what is really happening in my domain of conscious experience; when Sophie tells you that her pain is *bad*, she is also not capturing what it is fully like to be experiencing that pain. Along these lines, Kierkegaard made the profound suggestion that subjective life is essentially a *secret*:

> The entire essential content of subjective thought is essentially secret, because it cannot be directly communicated. This is the meaning of the secrecy… And the fact that this is the essential form of such truth, makes it impossible to express it in any other manner. (Kierkegaard 1846 [1941]: 73).

Similarly, Kant suggested that aesthetic experiences could not, at their core, be adequately understood or communicated through concepts or language:

> By an aesthetic idea, however, I mean that representation of the imagination that occasions much thinking though *without it being possible for any determinate thought, i.e. concept, to be adequate to it, which, consequently, no language fully attains or can make intelligible.* (Kant C3 5:314, emphasis added)

Kierkegaard and Kant are making the same essential point: that some of our most important experiences from within are fundamentally ineffable, taking on an internal meaning and character that cannot be reduced to language or description. Even the best of *normative* language cannot always fully communicate the complete nature of our internal experiences – this is why great poetry is admired, for it comes the closest to doing so, but can still only *approximate* a full communication of what has been experienced or understood.

When Sophie experiences her pain as bad, she is seeing a reason to get out of the pain – she is seeing the negative nature of her experience as *counting in favor* of her escaping the pain. To understand this point, one does not need years of meta-ethical reflection; one simply needs to feel what it is like to be in

genuine torment, for then one will understand that there is more to life than mere *description* about human psychology. Parfit's objection that there is no *normativity* on a mind-given theory can only succeed if he *equates* normativity with the idea that there are external or mind-independent normative standards. But that is not what it takes for there to be normativity, in the ordinary sense of normativity that we need to guide our lives with direction: for that, we simply need to see things as good and bad and beautiful and meaningful and valuable, from the driver's seat of being a human person. Normativity cannot be seen from the outside, but that does not mean it is not there.[126]

4.4 The Reality of Normativity

The final nihilistic concern is that, according to the mind-given theory, normativity is not in some important sense *real*. Murdoch, for instance, often emphasizes matters of reality, in her advancing a Platonism about the Good.[127] Similarly, Enoch has written that, in order to take normativity and morality *seriously*, we must adopt "an ontologically committed view" (2011: 7). Objections along these lines assume that, if normativity is a phenomenon of the mind, then it is in some sense *unreal*, in an objectionable manner – perhaps that it is an illusion, or a figment of the imagination. The concern would therefore be premised on two ideas: that for a theory of normativity to be satisfactory, it must capture that normativity is a part of *reality*; and that for something to be *real* in the relevant sense, it must be outside or independent of the mind.[128]

I believe that such an objection rests upon a mistaken view about the place of conscious experience in reality: it assumes that if something resides in

[126] Korsgaard makes this observation in a very poignant way: "Value, like freedom, is only directly accessible from within the standpoint of reflective consciousness… Trying to actually see the value of humanity from the third-person perspective is like trying to see the colours someone sees by cracking open his skull. From outside, all we can say is why he sees them" (1996: 124).
[127] See her remarks in *The Sovereignty of Good* (1970).
[128] I do not take recent *quietist* realists to be making this objection, who instead try to de-emphasize the metaphysical commitments of their accounts (e.g. Nagel 1986, Dworkin 2011, Parfit 2011, and Scanlon 2014).

the domain of phenomenology, then it is somehow less *real* than if it had an existence that was independent of the mind. But phenomenological facts are a part of *what there is*.[129] Phenomenological facts are a part of what there is, for the undeniable reason that there is *something it is like* to be experiencing: for instance, there is *something it is like* for you to be reading this text. We might say that this is one of the most *certain*, *foundational*, and *irrefutable* facts that a human being is capable of coming to ascertain. It is one of the most *certain* facts, because through *any* act of thought, we confirm the proposition that there is something it is like to have a mind. It is one of the most *foundational* facts, because having thoughts is the departure point for all further inquiry and reflection: it was how Descartes proceeded, when he began from the fact that he was *thinking* – a fact that resided within the phenomenological domain.[130] And it is one of the most *irrefutable* facts, because any conscious act of attempting to deny it, through thought, would only serve as verification for the point, since we would have then encountered a phenomenological fact. Thomas Nagel points out, correctly, that such facts are a part of *reality*:

> The subjectivity of consciousness is an irreducible feature of reality – without which we couldn't do physics or anything else – and it must occupy as fundamental a place in any credible world view as matter, energy, space, time, and numbers. (Nagel 1986: 7-8)

Street makes the same point:

> Existing independently of a subject's point of view on the world is not the only way of being "real." Normativity depends on the attitudes of beings who take

[129] I use the term 'phenomenological facts' to refer to facts about conscious experience: for instance, the fact that you now had an experience of reading this word is a phenomenological fact. We could also call them facts about 'phenomenal' consciousness (see §1.3).
[130] As Descartes writes: "*I am thinking therefore I exist*, was so secure and certain that it could not be shaken by any of the most extravagant suppositions of the sceptics, I judged that I could accept it without scruple, as the first principle of the philosophy I was seeking." (*A Discourse on the Method*, Part Four: 32, trans. MacLean).

> things to matter, but that doesn't mean it doesn't *exist*. (Street 2017: 121, emphasis added).

For this point to be true, we do not need to take a specific stance in the philosophy of mind on how phenomenological facts relate to physical facts, including how phenomenological facts might supervene on or reduce to physical facts – for even on the most *reductive* of views, there is still no denying that there is something *it is like* to be a human person, and that fact would then constitute an aspect of what there is.[131] Even for the most reductive theorist, there is still something it is like to *be* that theorist.

Since phenomenological facts are a part of reality, so is the normativity that can be experienced through our practical standpoints: there is nothing unreal about the fact that Sophie is in pain, nor is there anything unreal about the moment in which you find the Mahler Symphony or natural landscape to be beautiful. In the moment of experiencing pain or beauty, and the irreducible normativity of those experiences, you are having an experience with normativity that is just as *real* as if there were mind-independent normative truths about the matter. Perhaps this is what Hume had in mind when he indicated that we could experience "a *real* pain and pleasure" (T. 2.1.7.3), or when Nagel writes that "we and our personal perspectives *belong to the world*" (1986: 6, emphasis added). The fact that there is consciousness – that there is something it is like to be ourselves – is one of the most extraordinary facts about the world that we find ourselves in. What is even more extraordinary than that is that these conscious experiences can contain encounters with *normativity* – with things that fill us with meaning and delight and wonder. Whenever we have such an experience, we encounter a part of reality.

Conclusion

[131] Nagel as well, after the passage I quoted above, indicates that his position there does not imply a particular view yet in the philosophy of mind: "The ways in which mental phenomena are related to the brain… are matters that cannot now be settled, but the possibilities are appropriate subjects for philosophical speculation" (1986: 8).

This Chapter has argued that a mind-given theory of normativity need not be nihilistic. To the concern that sunsets will not always be beautiful, we can respond that normative potentialities are objective and unchanging, and that this can provide an anchor for our ordinary conception of value, as well as a sense of existential solace and appreciation for what our universe is like. To the concern that there are only descriptive facts about human psychology, we can respond that there is irreducible normativity from within the practical standpoint. And to the concern that normativity is not real, we can respond that phenomenological facts are just as much a part of reality as anything else. Through some creative uses of its resources, the mind-given theory can therefore achieve a picture of normativity that is deep and exciting.

5 Moral Objectivity

5.1 Introduction

Thus far, we have developed our mind-given theory in a number of ways: we have argued that it should focus on normative endorsements in general, rather than specific attitudes (Chapter 2); that there are a number of objective normative principles in the pre-moral domain (Chapter 3); and that we can achieve a picture of potentialities that mirrors the realist's mind-independent values (Chapter 4). These points have been intended to render our mind-given theory plausible and attractive as a holistic account of normativity.

But we have not yet discussed the most important topic concerning normativity: morality. Morality has always been the culminating aim of our discussion, for it is morality that gives us our most vivid instincts about normative objectivity. Many of my readers will be familiar with the fact that there is disagreement between different societies – as well as within societies – as to what the correct moral norms are. This is a well-known anthropological and sociological observation that we should not dispute. But what we must point out is that this observation alone is not enough to extinguish a deeper philosophical yearning for a vindication of moral objectivity. For as I indicated in the opening pages of this work, it is when we imagine a world in which Hitler won the war – and thereafter convinced a strong majority of people that his genocide was correct – that we will find ourselves with the following thought: that, *even in that world*, his genocide would never have been right, but would have remained profoundly *wrong*. So long as we do not give up on that thought, we will be on the quest for a vindication of moral objectivity: for we will need to show, on philosophical grounds, how there can be moral truths which obtain independently of the particular judgments that individuals or collectives may come to take.

Theism attempts to vindicate moral objectivity through an appeal to God, while realism asserts that there are brute and mind-independent moral facts. As

I explained earlier, I do not wish to reject either theism or realism, but rather, have found reasons to explore whether moral objectivity can be vindicated without relying upon these weightier assumptions (§1.2). So, instead of maintaining that morality is anchored in something that is *beyond* a human standpoint (theism and realism), and instead of maintaining that there is *no* objective anchor for morality at all (relativism), we want to see whether objective moral conclusions are entailed by the existing commitments of every human standpoint.[132]

This brings us to the question of whether we can find that there are objective moral truths that are *constitutive* of the human practical standpoint. This Chapter argues that we can. In particular, it argues that there is a moral principle, constitutive of the human practical standpoint, which demands that we treat every valuing perspective as generating normative reasons for ourselves, at least presumptively – i.e. that the valuing of every perspective *matters*, presumptively.[133] More broadly, I argue that this principle is the constitutive conclusion of any *intelligent* practical standpoint that shares certain core features with the human standpoint. I call this conclusion the **Principle of Universal Love.**[134] It maintains that the valuing of every perspective is as a source of presumptive normative reasons for ourselves. The principle therefore sets out an expansive moral standard governing human action, holding that we make a normative mistake whenever we turn a blind eye to the valuing of other beings in the world. I say *presumptively*, because not all valuing will turn out to be internally compatible with the Principle of Universal Love, and so some valuing

[132] This form of constructivist objectivity, or constitutivist objectivity, is sometimes said to be the original ambition of Kant, and a broadly similar strategy is pursued in a number of important recent works, which I examine in Chapters 6 to 10. Note again that this conception of objectivity here – as involving standards that are independent of our *particular* thoughts and attitudes, and of which we can be mistaken – is different from another sense of objectivity that concerns complete metaphysical independence of the mind (§1.3).

[133] I use the term 'valuing' in an expansive sense in this Chapter, as including many different normative attitudes. This is to be contrasted with how we used the term in §2.3, as referring to a specific and complex attitude that involves emotional aspects (as Scheffler 2011 and Tiberius 2018 understand the term). See also §5.2.

[134] This could also be called the Principle of Agape, insofar as *agape* is taken to be a concept that refers to love for all valuing perspectives.

will be ruled out from consideration on internal grounds – a point that I will elaborate upon, when I outline the general picture of morality that emerges (§5.4). For now, here is the schematic point: if our arguments are successful, then we can show that the Principle of Universal Love is the constitutive conclusion of our practical standpoint, and one that therefore binds us from within our own standpoints – rather than coming from somewhere that is beyond any standpoint at all.[135]

Why is the Principle of Universal Love a norm that is constitutive of the human practical standpoint? I shall present two independent arguments for this conclusion. The first argument maintains that every act of valuing by an intelligent agent involves a constitutive commitment to the proposition that valuing *qua* valuing is a sufficient condition for the deservingness of love. According to this argument, our own valuing implicitly involves certain interpersonal appeals, and as a result, our own valuing constitutively entails the conclusion that valuing *in general* – i.e. valuing *qua* valuing – is sufficient for love. I call this the **Conceptual Argument**. I then offer a second argument, which maintains that this same core proposition – that valuing *qua* valuing is a sufficient condition for love – is a commitment of the human practical standpoint, on *empirical* grounds. I call this the **Empirical Argument**. These two arguments, though presented as independent of one another, are united in advancing the same Principle of Universal Love, as a moral principle that inescapably arises from the human practical standpoint.

Before embarking on these arguments, it is worth saying a few words as to why one might find the core idea here attractive. First, these arguments move us from the idea that our *own* acts of valuing matter, to the more general point that we must also be concerned about the valuing of others, through the proposition that valuing *qua* valuing is a sufficient condition for love. These arguments therefore rely upon a longstanding idea in ethical thought: that, through reflection upon our points of view, we can come to see the normative

[135] Think here again about Rawls, who suggested that even an *objective* point of view "must always be from somewhere" (1993: 116).

significance of others. Consider, for instance, the injunction that 'we ought to treat others as we would like to be treated ourselves' – a teaching that has existed, in some form, in nearly all religious and ethical traditions.[136] This injunction makes essential reference both to how we *ourselves* would like to be treated, and to how we ought to treat *others* – a connective insight that I shall suggest can be rendered into a set of meta-ethical arguments. Second, I think these arguments also have a relatively simple and straightforward nature to them, in comparison to some of the elaborate proposals about intrapersonal unity or idealized counterparts found in the neighboring literature.[137] The Conceptual and Empirical Arguments both begin from the nature of our experiences with goodness and badness – from what is involved in our ordinary valuing – and extract a conclusion about what we hold valuing, *as valuing*, to deserve. This, I shall suggest, keeps us closer to our ordinary, phenomenological encounters with morality, in contrast to arguments that bring us into quite abstract and intricate philosophical terrain.

At the outset, let me also make a point about metaphor. In certain antirealist arguments for morality, it can sometimes seem that the theorist is seeking to bridge some *gap* through her argument: the gap between the individual and others, or to be more precise, from one's (non-moral) starting practical standpoint to the moral conclusion of the argument. This metaphor fits especially well when these arguments are formulated as making a conditional claim: *if* the agent is procedurally rational or logically consistent, *then* she must accept morality. For this can suggest that the argument is claiming that, *only after the standards of rationality or logic are applied*, does the agent *change* her existing commitments in order to conform with some new moral outlook. And this way of putting the argument invites the following important objection: what prior normative reason is there to be rational or logically consistent?[138]

---

[136] Often called the "Golden Rule," an injunction of this general form is found in Christianity, Judaism, Islam, Buddhism, Hinduism, Taoism, and Confucianism, among other traditions. As Simon Blackburn observes, this injunction can be "found in some form in almost every ethical tradition" (2001: 101).

[137] E.g. Korsgaard (2009) and Smith (2013), as discussed in Chapters 7 and 8.

[138] See Crisp (2006: 53), as well as Nozick (1981: 408) and Wallace (2012: 33).

I think that the mind-given theorist can benefit from shifting her metaphor: instead of accepting the idea that there is some gap to be bridged, the theorist should instead seek to show that *there is no gap in the first place*. In other words, what we should want to demonstrate is that *the very idea of a gap is an illusion*: that every human is *already* committed to morality, whether or not she realizes it. First, this is not only a shift in metaphorical framing, but has a rooting in the very nature of the arguments to be advanced: what we are claiming is not that procedural rationality or logical consistency must force the individual to adopt some *new* set of moral concerns, but rather that a significant moral conclusion *just is* what every individual is already committed to in her *existing* practical point of view. As such, there is nothing that the individual needs to change: it is merely about discerning what her practical standpoint *already* amounts to. Second, this shift can help the mind-given theorist better meet the challenge of how to establish the normativity of the moral conclusion, for it allows us to *circumvent* the objection mentioned above. It is not *in virtue* of following certain requirements of procedural rationality or logical consistency that an agent must then assent to morality: rather, she just is normatively committed to morality, as she already is.

Finally, note that, in making these arguments, I am not claiming that there are no further moral truths, beyond the Principle of Universal Love. There may be further moral truths, if theism or realism is true, or if any of the other constitutivist or constructivist arguments from the literature are successful. Rather than making an exhaustive claim about morality, what I want to suggest is that the Principle of Universal Love provides a *floor* for morality, such that even if set aside the weightier assumptions of theism and realism – and therefore start with the relativist's own metaphysical worldview – we can *still* show that humans are bound by a significant moral principle.

I begin by developing the Conceptual Argument (§5.2), before turning to the Empirical Argument (§5.3). After making these arguments, I then outline the general picture of morality that emerges, alongside addressing some substantive questions (§5.4). I also want to note that, because these arguments

are new and ambitious, my primary aim will be to *set out* these arguments as potential lines of thought, rather than to claim that they are invulnerable to further challenge. Indeed, what I explicitly hope to do is to put forward a set of premises – premises which I shall suggest are plausible and promising – that then *invite* further challenge and discussion, so as to carve out a new space for meta-ethical exploration.

5.2 The Conceptual Argument

The Conceptual Argument maintains that it is a constitutive commitment of the human practical standpoint – or more broadly, what I shall call an *intelligent* practical standpoint – that valuing *qua* valuing is a sufficient condition for the deservingness of love. Let me now elucidate what I have in mind with this statement.

We can begin with the idea that valuing *qua* valuing is a sufficient condition for the deservingness of love. Here, I use the term 'valuing' in an expansive sense, to include any kind of normative attachment in the world, ranging from rudimentary desires and pleasures and pains, to more sophisticated judgments and choices.[139] When we say that valuing *qua valuing* is a sufficient condition for the deservingness of love, we mean that valuing *as valuing* – i.e. the fact of valuing, *alone*, independent of other factors – is sufficient to deserve love. In other words, it would not be valuing *plus some other factor* that makes that valuing deserving of love; it would be the mere fact of valuing that triggers the sufficient condition. For instance, this would exclude the idea that it is *my* valuing that is specifically deserving of love, or the valuing *of persons with a certain race* that specifically deserves love, or even that it is valuing *of some particular species* that deserves love. It would be the principle that valuing, *no matter where it comes from*, is deserving of love. That would simply be the Principle of Universal Love, which maintains that the valuing of every standpoint is deserving of love.

[139] See Street (2016a: 176-177) for a similar usage.

Next, what do I mean by *love*? It is notable that this word – one that is so central to our ordinary life – is not often found in recent ethical discussion.[140] As I shall use the term, to *love* the valuing of another being is to treat that valuing *as a source of normative reasons for oneself*. It is to see that being's needs or concerns as having normative import for one's own deliberations and actions – it is to see that being's needs and concerns as *mattering*. This framing is admittedly general, but as I shall explain, this is in part a feature of the view, with substantive implications that we shall explicate (§5.4). There are many types of normative reasons that might be fitting of love: for instance, there are familiar reasons to promote, respect, or not interfere with, but there can also be reasons to empathize with, to care for, to cherish, to protect, to remember, to honor, or even just to think about or see something in a certain way. The point here is that, if we take another's valuing to be deserving of love, then it cannot be something that bears *no* normative weight upon our deliberations or actions – it is not something that simply can be ignored, *as a reason*. I therefore use the term *love* both as a philosophical term of art, and in a way that tracks some of its ordinary meaning.

Before turning to the argument proper, let me also discuss the kind of practical standpoint that is meant to constitutively entail the Principle of Universal Love. I have mentioned the *human* practical standpoint, since I am interested in saying something that applies to you and me; but more broadly, we could speak of the *intelligent* practical standpoint. I have chosen to use the term 'intelligent,' rather than the term 'rational,' because the latter is one that comes with greater philosophical baggage. For instance, different authors use 'rational' in different ways, and it often comes in a *normatively loaded* sense, in which the notion of rationality already involves certain embedded normative standards.[141] Unlike certain accounts of rationality, I define a being as *intelligent* – as having an intelligent practical standpoint – insofar as she has certain *descriptive*

[140] We can recall Murdoch: "We need a moral philosophy in which the concept of love, so rarely mentioned now by philosophers, can once again be made central" (1970: 45).
[141] This is one way in which an implicit form of realism can be found in some accounts, as I argue in Chapter 9 about Markovits.

capacities: if she is capable of deliberating and acting upon normative reasons; if she is capable of seeing herself as existing as one among other intelligent minds, in which those other minds are understood as capable of influencing the course of her own life; and if she is capable of seeing that all other practical standpoints have normative attachments of their own, i.e. that other creatures also value in the broadest sense.[142] These capacities describe certain kinds of beings, but they do not come embedded with normative standards. Now, human beings, at least in their focal case, have intelligent practical standpoints in this sense – where my sense of intelligence is broad, and involves certain minimal forms of relational and emotional intelligence, in the capacities of awareness about others.[143] It is furthermore an interesting question whether there are other forms of life that have intelligent practical standpoints: insofar as there are, these forms of life would also be constitutively committed to the Principle of Universal Love, according to the Conceptual Argument. I shall refer to any being that has an intelligent practical standpoint as an *intelligent agent*.

The question is why we should think that it is constitutive of the intelligent practical standpoint that one must maintain that valuing *qua* valuing is a sufficient condition for the deservingness of love. The **Conceptual Argument** seeks to establish this conclusion, through four premises:

1. Every intelligent agent *values*, in the broadest sense, as an *analytic* matter: every intelligent agent has at least some kinds of ends or normative attachments in the world. (*Valuing Premise*)

[142] Some of these points invoke echoes of Nagel (1970), though the arguments we use them for are quite distinct.

[143] I therefore do not use the term 'intelligent' with any of the connotations it sometimes has, of being 'smart' or 'having a high IQ': I am appealing to a very broad sense of intelligence that includes practically all humans. However, there are some marginal cases in which a human is not capable of deliberating or acting upon reasons – for instance, those in permanent vegetative states – who will be excluded as subjects of moral requirements, though these humans will still be those who deserve love, as *recipients* of the Principle of Universal Love. Similarly, a creature who lacks the ability to understand that others are real – but who can only see them as objects, like the rocks – will also not be a subject of moral requirements, but will be a recipient of the Principle of Universal Love, insofar as the creature values.

2. If an intelligent agent values something, then she makes the implicit and constitutive demand on all other possible intelligent agents to treat her valuing *as a reason*, i.e. to treat her valuing as deserving of love. (*Interpersonal Premise*)

3. If an intelligent agent demands that all other possible intelligent agents treat her valuing *as a reason*, then she constitutively seeks to *reach* those intelligent agents with her claim through an act of reasoning, which entails that she must formulate her claim in a way that is understandable and applicable to them. (*Reasoning Premise*)

4. If an intelligent agent must maintain her claim in a way that is understandable and applicable to all other possible intelligent agents, then she must commit to the principle that valuing *qua* valuing is a sufficient condition for the deservingness of love. (*Sufficiency Premise*)

5. Conclusion: Every intelligent agent is constitutively committed to the principle that valuing *qua* valuing is a sufficient condition for the deservingness of love.

The Conceptual Argument moves from a relatively uncontroversial idea – the idea that every intelligent agent *values*, in the broadest sense – to the more expansive conclusion that every intelligent agent is beholden to a moral principle that valuing *qua* valuing is a sufficient condition for love. Let me now explicate each of these premises.

We begin with the first premise, the *Valuing Premise*, which states that every intelligent values, in the broadest sense, as an analytic matter. This premise is meant to be the most straightforward one. I have said that the target of our argument is the class of intelligent agents, which by definition, have *practical* standpoints – standpoints that involve at least some minimal layer of valence, or normative experience, to them. Therefore, as a definitional matter, all intelligent agents have certain normative attachments, in the broadest sense.

So, the *Valuing Premise* essentially makes an analytic observation about the kinds of beings that are the target of our argument.[144]

Next, we move to a more significant premise, the *Interpersonal Premise*, which states that if an intelligent agent values something, then she makes an implicit and constitutive demand on all other possible intelligent agents to treat her valuing *as a reason*, i.e. to treat her valuing as deserving of love. The idea here is that, if one is not making such a demand on all other possible intelligent agents, then one is not *valuing at all, as an intelligent agent*. Remember that I defined the intelligent agent as one who is conceptually aware of the possible impacts of others' actions on the success or failure of her own ends. Because we have this conceptual awareness of how others can at least potentially understand and impact our own ends, our setting a value on something implicitly involves making a claim on others as well. For instance, imagine that you are being tortured, but there are some others who might relieve you of your torture – others who might *understand* your torture as bad, and then *do* something about it. To *value* not being tortured is to implicitly ask them not to torture you: more specifically, it is to ask them to see that they have some reason to relieve you of your suffering.[145] If you maintained that they had no reason to help you at all – that it makes no difference whether they see your pain as mattering – then you wouldn't be *valuing* not being tortured, as an intelligent agent. Even the enjoyment of an apple involves the implicit demand, from one intelligent agent to another, not to swipe away the apple from one's hand. We might say that

[144] Insofar as some creature did not have any normative attachments at all, they would also not have a practical standpoint, and would therefore not be the type of creature for whom we would intuitively want to vindicate as a subject of moral requirements.

[145] One could object that, in valuing something, all we commit to is *wanting* someone else not to interfere, but that this does not amount to a claim on their *reasons*. Williams seems to make such an objection (1985: 68-69), when discussing an argument that seems to point to the important work of Alan Gewirth (1978). The response to Williams is that, if an *intelligent agent* values, but attempts to deny that she is also making a claim on others' reasons, then she is simply displaying a certain lack of self-knowledge or conceptual knowledge as an intelligent agent. For the idea is that *intelligent* valuing involves this implicit claim on others, in virtue of its conceptual awareness of the nature and capacities of other intelligent minds; perhaps for some less conceptually aware minds, e.g. for certain non-human animals, valuing may not come along with this implicit claim on others. But when one is aware that other intelligent minds are capable of actually *understanding* and *seeing* one's own plight as a reason for them, genuinely *valuing* something comes along with making an implicit *claim* on them.

intelligent valuing – valuing that is done with a conceptual awareness of others' ability to understand and affect or valuing – involves making these implicit claims on others.

Notice that this implicit and constitutive demand is made towards *all other possible intelligent minds*. This point is important. First, the demand is made towards other *intelligent minds*, because the distinguishing mark of our own interpersonal intelligence is that we recognize that it is *these* creatures in the world that could actually understand and respond to the kind of claim that we are making on them. Unlike the wind or a mosquito, other intelligent minds are at least potentially *responsive* to the claim we are implicitly making on them: if the wind causes a branch to swipe at us, we might try to move out of the way, or if the mosquito is hovering around us, we might try to swat it away, but we would not be making a *claim* upon the wind or mosquito, for we are aware that no such claim could ever reach them. We would not be asking for potential *love* from them, but merely treating them as obstacles to be navigated around. Second, it is also important that the *Interpersonal Premise* is formulated as concerning *all possible* intelligent minds. This is because, at bottom, the constitutive claim is that *if* there were some intelligent creature around, then that creature has reason to see that our valuing matters. When we walk around in societies that are surrounded by people, the interpersonal aspect of our valuing is easy to see; but I want to suggest that the point holds even on a remote island as well. When you fall into a deep hole on that island, and mourn what looks like the impending loss of your own life, you are making the claim that *insofar as there is anyone around* – "can anyone hear me?" – that person would have a reason to throw you a rope and respond to fact that you are about to lose your life. Indeed, let me suggest that the deepest form of this conditional demand happens even if we *know* that there is no such living creature around us: for it can still take the form of a prayer or hope that, *insofar as there is a God* – an intelligent mind who *might* be able to intervene and help – then our suffering is something that ought to be treated as normatively significant, as mattering. Indeed, this is why many humans in stranded or otherwise trying circumstances – when it seems that no

else can help with a dying loved one, or some other profound plight – do turn to implicit forms of prayer or hope. That is evidence of the Interpersonal Premise in action: in such cases, there is the implicit claim being made on possible intelligent minds to see our valuing as deserving of love.

This brings us to the third premise, the *Reasoning Premise*, which maintains that when we make these implicit and constitutive demands on all other possible intelligent minds, we are attempting to *reach* those minds through an act of reasoning, and that means that we must formulate our claim in a way that is understandable and applicable to all of them. As I mentioned earlier, we do not make these constitutive demands on the wind or a mosquito – because we ourselves are intelligent minds, with an awareness of how other intelligent minds could in principle understand and respond to our claims on them, our claims specifically have *those other intelligent minds* as their constitutive focus. And the mark of that focus is that we are implicitly trying to *reason with them* – we are trying to *reach* them with our claim, to make them see or understand something about our plight that would ground a reason for them. That is why our claims on them must be formulated in a way that is intelligible to them, and which they cannot easily reject as irrelevant to them, for only then would we be making our best attempt at actually *reaching* them. Otherwise, it would simply be like we were dodging the wind or the mosquito, rather than making a *claim* on someone who we understand is another intelligent person, capable of thinking through and responding to our claim. So, bring yourself back to the torture cell, where the others are beating you: as you are suffering, and therefore making implicit demands on them to stop, you are constitutively seeking to tell them something about your suffering that they can understand and respond to – to say something that they can find intelligible and accept as having normative implications for them, at least as an *attempt* to get them to see that your pain is something *worth* being relieved from. We might therefore say that intelligent valuing *beckons towards an interpersonal understanding*.

And that brings us to our fourth premise: the *Sufficiency Premise* maintains that, if an intelligent agent must formulate her claim in a way that is

understandable and applicable to all other possible intelligent agents, then she must commit to the principle that valuing *qua* valuing is a sufficient condition for the deservingness of love. Why is this so? It is because all other alternative claims would not constitute the most intelligible and widely relevant attempts at reaching all those who could possibly affect us: for any alternative attempt might always bear no direct relevance to our potential interlocutor, as well as involve no intelligible explanatory connection to them. For instance, imagine presenting a claim to another intelligent agent that leaves them entirely cold, in that it appeals to some necessary condition for love that has no bearing or relevance to them at all. Someone of a different race is torturing you, and you tell them: "it is only *because I am of my race* – rather than *yours* – that you should stop torturing me!" Not only is it difficult for them to find some *explanatory* connection between your race and demand for relief, but moreover, you have made a claim that leaves them entirely cold – one that has no relevance to them. Or, imagine that intelligent aliens have descended upon Earth, and are moments away from eradicating all that we care about. If we told these aliens: "it is only because *we are the particular biological species we are* – rather than *your biological species* – that our valuing deserves to be respected by you," we would have made a claim that is far from our *best* attempt to reason with them. Whenever we attempt to add some contingent, necessary condition to our claim for love, we make a claim that potentially has no bearing on some other possible intelligent mind. Even if we formulated our claim on other intelligent minds in terms of the *particular substantive content* of our valuing, we would potentially leave them cold, since they could be completely indifferent to that substantive content: humans might care about their mountains and lakes, at a substantive level, but the aliens might simply see those mountains and lakes as terrain to be navigated around.

It is only when we formulate our claim as valuing *qua* valuing deserving love that we appeal to something that all other possible intelligent agents can understand, find relevant to them, and see some explanatory connection in. When someone is torturing you, and you tell them that your valuing is just like

theirs – it is valuing *qua* valuing that is meritorious of love – you make a claim that is fully intelligible and relevant to them, for all possible intelligent agents, by definition, *value* in this way as well. And because I have argued that our valuing makes an implicit claim upon *all other possible* intelligent minds, the only claim we could make that genuinely attempts to reach them *all* appeals to a condition that is necessarily shared between all intelligent minds: *that there is valuing*, i.e. valuing *qua* valuing. When we formulate our demand in these general terms, rather than with some other necessary condition, we appeal to an idea that all other possible intelligent minds can find resonance, intelligibility, and relevance in. Moreover, unlike other alternative formulations of a claim that are more contingent in nature, this wide claim is one that involves a clearer explanatory connection: for those other intelligent minds, who also know what valuing is like, can at least see *why* valuing would lead to the demand that that valuing is deserving of love from them, rather than if some particular race or specieshood were appealed to. And recall that these points continue to hold, even if we were on a remote island: for even when we value on that island, we would still be implicitly seeking to reach all other *possible* intelligent minds, and it would therefore still only be valuing *qua* valuing that could reach all of them, no matter what the other differences.

That brings our argument to a conclusion: from the idea that intelligent agents *value at all*, we reach the constitutive result that intelligent agents must maintain that valuing *qua* valuing is a sufficient condition for the deservingness of love. For because we make implicit claims on others that must contain this sufficient condition – claims on others that we imbue *normativity* into, from our own acts of valuing – we therefore render normative for ourselves a principle that involves universal love for others. For if valuing *qua* valuing is a sufficient condition for the deservingness of love, then this entails that valuing, *anywhere* – independent of class, identity, status, race, gender, specieshood, or substantive content – is presumptively deserving of love.[146] We therefore imbue normativity

[146] Again, I will soon explain why I say *presumptively* here (§5.4).

into a principle that demands that we see all valuing as presumptively mattering. That is the Principle of Universal Love.

Let me make a few closing comments on this argument. First, as will be evident to my reader, there is some appeal to our *social nature* in this argument: intelligent minds are conceived of as ones that display a certain degree of sensitivity to the fact that other intelligent minds can understand our claims on them, and in virtue of that understanding, to respond with a loving influence upon our lives.[147] However, this appeal to our social nature is not a *directly* normative one, in that it does not assume that all intelligent agents already care about all others. If we imagine Caligula as someone who is an intelligent agent in our sense, in that he can understand *that* others are suffering, but professes not to care, we can still display the force of the argument to him: for insofar as Caligula cares about something – such as his own not being tortured – then he will be making the constitutive commitments that we have outlined. Caligula, though he is not much of a 'social person' in the standard sense, is committed to the Principle of Universal Love, insofar as he *values* with a conceptual awareness of how others could understand and affect his ends, for then he would need to formulate his implicit claims under the principle of valuing *qua* valuing as a sufficient condition for love, as the argument goes. Second, let me note that, although certain points of this argument might sound empirical at times, the points are made in the space of conceptual relations and constitutive entailments: it is not that intelligent agents must actually make these claims out loud on others, but rather that the nature of their valuing involves a constitutive commitment to the kinds of claims on others that I have described. And again, the mark of a constitutive claim here is that, in already being committed to X, one just is committed to Y as well, such that if one denies that one is committed to Y, then one will also be denying that one is committed to X (but which one cannot do, insofar as it is a secure premise that one is already committed to X). So, if an

[147] One can also compare what some Kantians say about how reasoning involves intelligibility or justifiability to others (O'Neill 1996: 51-59, 2000: 24, and 2003: 538; Markovits 2014: 134; Markovits & Walden 2021: 328-333), though I do not describe my own account here as Kantian.

intelligent mind *values* ('X'), that mind just is committed to the idea that valuing *qua* valuing is sufficient for love ('Y'), such that if one attempts to deny the latter idea, one would then just not be *valuing* at all – but all intelligent agents *do* value, which is what makes the commitment inescapable for intelligent agents.

Finally, where does the *normativity* enter into this argument? Unlike certain conceptions of rationality, we have not directly built normativity into our conception of intelligence; rather, normativity enters into the argument when we bestow normativity into the world, through our own starting acts of valuing. Following the constitutive relations I have described, we then render normative a series of claims that result in the principle that valuing *qua* valuing is a sufficient condition for love – where this is a principle that we have constitutively made normative for ourselves, through our taking things to matter in the first place. It is therefore worth emphasizing again that, in the design of this argument, we can circumvent the question of why the conclusion is normative for the agent: for I have claimed that this moral principle *just is* the constitutive commitment of our ordinary acts of rendering the world into a normative place. Morality, we might say, is constitutive of intelligent life.

5.3 The Empirical Argument

The Conceptual Argument draws a series of constitutive relations that begin from any act of intelligent valuing, leading to the Principle of Universal Love. Let me now suggest that there are some interesting reasons to explore whether a similar argument can be made on *empirical* grounds. This brings us to the **Empirical Argument**.

The Empirical Argument asserts that human beings are committed to the proposition that valuing *qua* valuing is a sufficient condition for love, *as a matter of empirical fact*. The argument therefore does not try to draw a constitutivist line of connection from any act of intelligent valuing to a moral conclusion, but simply asserts the core idea about valuing *qua* valuing as an actual, dominant

commitment of the human practical standpoint.[148] This is not the claim that all humans will profess this principle out loud, or even the claim that they always necessarily recognize it in a self-conscious manner. But it is the claim that, as a matter of fact, we commit ourselves to this principle when we value, such that it serves as an authoritative, mind-given normative principle for human life.

The heart of the Empirical Argument therefore rests upon an empirical assertion – one which I cannot prove from the armchair, and which *could* turn out to be false. What reason is there, then, to find the Empirical Argument plausible or attractive? I shall suggest that there is an interesting dialectical route to be pursued here: we can argue that ordinary reflection renders the empirical assertion a plausible one, and that some further philosophical observations can show that the empirical assertion is quite difficult to falsify, and which allows it to be insisted upon as true, through a noble form of dogmatism.

Imagine that you are being tortured by a throng of people – your suffering is intense and excruciating. The Empirical Argument makes the *empirical assertion* that, when a human being undergoes such a normative experience, she *sees* or *understands* that it is the *nature of the pain* – or the *nature of valuing*, more generally – that grounds its deservingness of love. In other words, as the lashes are coming down on you, my empirical assertion is that you would not genuinely think to yourself that 'it is only *because I am a man*, or *because I have these particular characteristics*, that I am *deserving* to be relieved from this intense suffering.' My empirical assertion is that, in such a moment of having a normative experience, you would understand that it is the *nature of the pain, alone* – independent of these other factors – which is *enough* for that pain to be deserving of relief. The Empirical Argument therefore makes a postulate about what human beings do indeed commit themselves to, as an operative normative principle, when undergoing normative experiences: they see that it is the pain *qua* pain, or valuing *qua* valuing, that is a sufficient

[148] In this way, the argument has a more distinctive focus on the human being as a subject of empirical assertion, rather than a stronger conceptual claim to be made about *all intelligent beings, as intelligent beings* – though the empirical assertion I make still might apply to them.

condition for its deservingness of love. So, for instance, when it is often heard that people in very difficult circumstances simply 'wish to be treated like a human being,' we can see these individuals relinquishing an appeal to particular aspects of their identity as grounding a deservingness of love, and instead appealing to the thing that they have in common with all other humans – i.e. the ability to understand the nature of pain and other normative experiences, like the experience of needing help. The empirical assertion I am putting forward is that, in cases of pain or valuing more generally, we come to see, as a matter of our actual, internal commitments, that valuing *qua* valuing is a sufficient condition for love. Perhaps in the way that, if some human knows a language, she cannot help but understand the words that she is hearing, the idea here is that, if some human really undergoes a normative experience, she cannot help but understand that it is valuing *itself* that is deserving of love.

Now, I have not claimed that I can prove that this empirical assertion is true, and again, this assertion *could* turn out to be false. Moreover, anyone could contest this idea out loud, whenever they claim that it is not valuing *qua* valuing that is a sufficient condition for love. A King might exclaim that it is his Kinghood that makes his valuing specifically deserving of love; a racist might exclaim that it is her race that makes her valuing specifically deserving of love; and an egoist might exclaim that it is his being *him* that makes his valuing specifically deserving of love. To these apparent counterexamples, I now want to make three dialectical moves, so as to embed the empirical assertion I have made in a relatively unique and protected philosophical space.

First, notice that even when we are met with these apparent counterexamples, it is hard to *falsify* the empirical assertion, simply on the basis of having observed these exclamations from an external standpoint. For the empirical assertion is *not* one that concerns what people state or report on the surface, but about what human beings in fact internally commit to, when undergoing normative experience. Because there is a general problem of accessing, with certainty, the internal contents of other minds, it is at least possible that, even when people make the aforementioned claims aloud, on the

inside, they do in fact recognize that the nature of pain or valuing is itself sufficient for love. In fact, a rather deep point emerges: the only way in which the empirical assertion might be demonstrated to be false is if we, ourselves, could know *with certainty*, in moments of great pain or valuing, that we do not assent to the sufficient condition. So even in principle, only my reader has the potential to ultimately know herself whether the empirical assertion can be successfully falsified or not. To the aforementioned counterexamples then, where for instance the King says certain things out loud about his own Kinghood, we can always leave open the possibility that such an individual is not honestly or accurately reporting the actual, dominant commitment of his practical standpoint. The King might claim that it is being the King that makes his pain deserving of love; but when he is pain, we may insist that he can *see* that it is the pain itself that is deserving of love, and therefore that he renders normative for himself that idea, on the basis of his acknowledgment of it from within his practical standpoint.

Second, let me note that, even introspectively, a person might be mistaken about her own deepest attitudinal commitments at times, on the assumption that humans are not always perfectly internally luminous. For instance, sometimes we have a certain attitude towards someone, even when don't consciously realize it; other times, we might see the meaning or point in something, even before that becomes a self-conscious thought. In general, as Williams notes, someone can be "ignorant" of some attitude that they have, though that attitude would still ground an internal, normative reason (1981: 103). Korsgaard makes this point against perfect luminosity as well:

> Some philosophers have supposed… that our minds are somehow internally luminous, and that their contents are completely accessible to us… and so that introspection yields certain knowledge of the self. Like Kant, and many philosophers nowadays, I do not think that this is true. Our knowledge of our own mental states and activities is no more certain than anything else. (Korsgaard 1996: 62)

This point creates space for the idea that, even if we do not profess to ourselves, in our own internal language or explicitly recognized patterns of thought, that valuing *qua* valuing is sufficient for love, we might still be mistaken about the fact that we *do* assent to that attitudinal commitment, in undergoing normative experience. For, the thought that motivated our empirical assertion was that, as part of experiencing the nature of pain or valuing, humans come to understand that it is the pain or valuing itself that is deserving of love; and here, it might be that the nature of pain gives us a certain kind of background understanding or commitment to the sufficient condition at hand, even when do not realize it. So, even if we ourselves seem not to assent to the empirical assertion that I have made, that point may not fully secure a falsification of the claim, either.

Third, given that the empirical assertion is hard to falsify, let me suggest that there can be a noble form of dogmatism in insisting upon its truth, even though we cannot *prove* that it is true. Remember that we are seeking to achieve a vindication of moral objectivity – a task that is philosophical important, in virtue of our strong pre-theoretic intuitions about the matter. Moreover, our project is set in contrast to *realism*, which makes dogmatic claims about a domain of brute and mind-independent normative facts. The dialectical point here is that our form of dogmatism might be *comparatively benign*, in comparison to the kind of dogmatism that the realist must engage in. First, the realist must make assertions about something that raises heavier *metaphysical* questions: because the realist makes claims about mind-independent normative facts – facts which concern something more abstract and less tangible than the actual commitments of the human rights before us – realism can engender more metaphysical concerns about what these mind-independent normative facts *are*. At least at the metaphysical level, our empirical assertion about what human beings are actually committed to – which is *metaphysically* on a par with a claim like 'all human beings, as a matter of empirical fact, want to be happy' – avoids raising the same kind of metaphysical concerns that Mackie had about realism (1977: 38), not to mention that our assertion is more metaphysically parsimonious. Second, consider that realism also faces more *explanatory*

difficulties than our assertion, for the realist must assert that brute, fundamental normative facts *have no further explanation*. In contrast, our assertion about human beings is in principle open to further explanation of a psychological, biological, evolutionary, historical, or even theological nature. And finally, because we have made an assertion about a *descriptive* claim, about what human beings are committed to, rather than a directly *normative* claim, our empirical assertion also seems to fare better in avoiding certain practical and motivational objections that are leveled at realism. So, here is the point: if we must be dogmatic at all, there may be some attraction to *this* form of dogmatism as more philosophically benign and acceptable than the weightier form of dogmatism that the realist might be taken to need.

That is the Empirical Argument. It makes an empirical assertion, which we can suggest that there are some reasons to find plausible, then embeds that empirical assertion in a series of dialectical observations that makes it hard to falsify, and finally insists that there can be a noble form of dogmatism in insisting upon it as true. Like realism, the Empirical Argument is to some extent dogmatic, but it roots itself in a dogmatism about something much less ethereal. It makes the claim that, when we are facing pain or torture or other forms of normative experience, we would not deep down think that *it is because I am me, or because I am a King, or because I am of this race*, that my pain or torture or normative experience is deserving of love from others. We would see and understand the truth that the pain alone – or that the normativity of the experience, alone – is what grounds a deservingness of love. And that would be to empirically commit ourselves to the operative normative principle that valuing *qua* valuing is a sufficient condition for love, i.e. that the valuing of every perspective is deserving of love. The Empirical Argument therefore posits that human beings are distinctively shaped for goodness – not in a straightforward manner, and not in a way that we all actively realize – but in a way that involves our actual normative commitment, through a simple dimension of our own experiences, to an expansive standard of love towards others.

As a closing point, let me observe that some of these dialectical moves find an interesting parallel in the important meta-ethical work of Adrian Moore (2003). Moore's account of morality is one that sets aside realism, but which still seeks to be objective.[149] Roughly speaking, one of Moore's central claims is that, *insofar as we are rational*, then we must be committed to the normative implications of our using certain ethical concepts, such as of making a promise.[150] But the fundamental question then arises: what reason is there to be rational?[151] It is here that Moore makes a dialectical move that is similar to ours: he insists, as a contingent matter of fact, that humans are committed to being rational. Calling this the Basic Idea, he writes that "there is a nisus in all of us, more fundamental than any other, towards rationality" (2003: xviii). He makes clear that this is a contingent assertion that he cannot prove: "the Basic Idea, if it is true, is a deep contingent truth" (2003: 170), and "there is no way to establish the truth of the Basic Idea – even though I *hold* it to be true" (2003: 137, emphasis added). But Moore also makes clear that his assertion cannot be *disproven*: "I no more think the Basic Idea can be refuted than I think that it can be established" (2003: 170). And this creates space for Moore to insist upon it as true, as well to "hope that it is true – in the absence of any way of refuting it" (2003: 170).[152] There is something notable to the dialectical point that Moore is after, which comes with an echo of Kant: that if there is some important philosophical claim that needs vindication – one which we may even have practical reasons to vindicate – then, if we find that such a claim is significantly buffered from both proof and falsifiability, we may simply insist upon it as true, in a uniquely protected philosophical space.

5.4 The Demands of Love

[149] See Moore's discussion of "world-directed objectivity" (2003: 5-11), which we can understand as realism.
[150] See Moore (2003: 65).
[151] See Moore (2003: 29 & 130).
[152] Moore even adds that this point may come with an "aesthetic" or "religious" dimension to it (2003: 171).

I have just offered two independent arguments leading to the Principle of Universal Love, as an inescapable moral principle of the human practical standpoint. In this final section, we turn to address some questions about the picture of morality that emerges as a result.

At the most general level, the Principle of Universal Love has the injunction that every perspective is a valid source of normative reasons for intelligent minds – for minds like the human one, who can see that those perspectives have normative attachments of their own. We must therefore treat the concerns and needs and sufferings of other perspectives as ones that are presumptively important: no valuing perspective can simply be *ignored* as normatively irrelevant. We cannot treat the valuing of others as mere tools or obstacles to be navigated around: we must treat them *as presumptive reasons*. Hence, the Principle of Universal Love is expansive, finding its scope in all of the sentient and valuing perspectives in our world. And now, a number of questions can emerge about what sorts of substantive demands this principle leads to. And although this dissertation is not designed as a piece of normative ethics – it is primarily a *meta-ethical* work – it will be important to address at least some of these substantive questions.

First, throughout the arguments I have made, I have said that others' valuing are *presumptive* sources of reasons for us – and this is because, insofar as they are also committed to the Principle of Universal Love, there are certain ends or forms of valuing that will be internally contradictory with the idea that valuing *qua* valuing is a sufficient condition for love. For instance, if someone holds a first-order end that directly contradicts this idea – such as an end to torture someone else who wishes not to be tortured – then this person will have failed by the Principle of Universal Love, in even setting that as their end. Because the Principle of Universal Love rules out such ends as illegitimate, we can say that it only demands that we love those ends in others that do not, themselves, cause a contradiction with the principle. The injunction of the Principle is therefore to have love for the valuing of all others, where that valuing

is not *itself* prohibited by the same principle of having love for the valuing of all others.[153]

Second, what is the *strength* or *categoricity* of the moral conclusion that has been established: how demanding is the morality that emerges? In both arguments, I have claimed that the result is that valuing *qua* valuing is a sufficient condition for love – a principle that we reach from the nature of our *own* demands for love, on the basis of our *own* valuing. This therefore promises an elegant result: that we commit to the *same amount of love to others as we wish for ourselves, on the basis of our own valuing and implicit claims on others.* This answer is not a further substantive norm to be derived from additional argument, but rather lies at the nucleus of the arguments we have already presented. So, insofar as we have aims and ends that we constitutively demand that others see *as reasons*, we must bestow that same amount of reason-giving status to the valuing of others. For the sufficient condition just involves two sides of the same coin. There is therefore a very general equality that emerges in the moral result of these arguments.

Third, what does the principle tell us about the everyday business of making normative decisions, including in difficult choice circumstances? The Principle of Universal Love maintains that everyone's perspective matters, in that we must treat their considered valuing *as reasons*, but it is also a feature of the view that a certain amount of imprecision can be admitted to particular first-order questions, as Aristotle noted.[154] What I have hoped to demonstrate as a guiding principle is that everyone's perspective *matters*, in that we cannot simply ignore them; so for instance, if we choose to do something that incurs a harm on someone else, then we will be acting against at least one of our *pro tanto* reasons. But there may be more than one way in which we can show love

[153] Moreover, insofar as our own pre-moral reasons are determined by our own *counterfactual* endorsements (see §3.2), and it is towards those counterfactual endorsements that we demand love for, we might find that these arguments, amended with this point in mind, demonstrate that it is the counterfactual valuing of others that most deserves our love.

[154] As Aristotle writes: "the accounts we demand should be appropriate to their subject-matter; and the spheres of actions and of what is good for us, like those of health, have nothing fixed about them. Since the general account lacks precision, the account at the level of particulars is even less precise" (Nicomachean Ethics, 1104a, trans. Crisp).

towards others, in any particular circumstance, for there are so many perspectives that require love; and the idea that everyone's valuing matters may also take us beyond some of the straightforward first-order questions that often assume a single answer to matters of right and wrong in particular circumstances, to the broader and more expansive question of how everyone's perspective can be honored and seen *as a reason*.[155] Indeed, the Principle of Universal Love is not just about action, but also extends to thought: for we respect the principle when we see people passing on the street, or in the grocery store, and we *will* their good, thereby treat their valuing *as normatively significant*. More generally, the Principle of Universal Love may require us to work on our character or virtue in subtle and sophisticated ways – including changing our *normative lens* upon the world (§2.5) – to better fit towards being sensitive and loving towards the valuing of all, and to become more empathetic and naturally disposed towards generous action. Moreover, the principle also comes with implications for political life, for if one has the opportunity to act on a stage that impacts a large number of lives, one will still be under the injunction to treat everyone's perspective as mattering. And note that nothing about the picture that has been suggested thus far rules out the possibility of challenging dilemmas: sometimes, in the real world, it is not possible for one individual to perfectly promote or respect the ends of all others at once, given the available courses of action – but that is the nature of practical life, and the guiding injunction is that valuing, wherever it can be found, is *worthy* and *deserving* of love. Even when acting in these difficult circumstances, we must be committed to the conclusion that everyone's perspective matters, and to let that truth inform our deliberations.

Fourth, it is worth emphasizing that, insofar as valuing *qua* valuing is a sufficient condition for love, there arises obligations towards non-human animals: for even these animals have valenced practical standpoints, which

[155] For instance, in the infamous Trolley problem, the relevant normative and moral questions are not exhausted after we decide whether to turn the Trolley or not; for there are also questions that emerge about *attending to* the victims, such as honoring their lives and comforting their families, as well as how much emotion or sensitivity we show towards the fact that the tragic choice between saving different lives had to be made at all.

involve forms of valuing. Our conclusion was that valuing *qua* valuing is deserving of love – and since animals value as well, their valuing is deserving of love.

Finally, let us return to our motivating question: how has this given us a floor for moral objectivity? We have articulated a principle that is constitutive of the human practical standpoint – one that has normative force for us – which instructs us to see others' perspectives as mattering, just as we demand for ourselves. And notice that, although we normatively commit ourselves to the idea that valuing *qua* valuing is sufficient for love, we might often fail to realize what this idea fully entails – or we might simply fail to actually follow through on the principle, in practice. In other words, just because we commit ourselves to this principle, does not mean it is easy to recognize or live up to its implications. So, morality is objective, on the account we have offered, because when we are insensitive to the perspective of others – either in our own actions or thoughts – then we violate a moral principle that we cannot help but bestow. I have not claimed that there are no *further* moral standards, beyond the Principle of Universal Love: I have simply claimed that human beings cannot help but be committed to this principle, as an ethical *floor* that constitutively governs their normative lives.

Conclusion

This Chapter has sought to develop a single core idea: that our way of valuing involves a commitment to seeing the normative significance of others. This idea is at the heart of most religious and ethical traditions, and the suggestion at hand is that it can be developed into a distinctive set of meta-ethical arguments. The Conceptual Argument maintains that all intelligent agents must assent to a sufficient condition that results in the Principle of Universal Love, through a series of constitutive relations that follow from any act of intelligent valuing. The Empirical Argument asserts the same principle as one that we are committed to, on empirical grounds – and notes that there are some unique attractions to

how this claim can be asserted. In both arguments we reach the requirement that we must bestow love upon a wide swath of ends in others, as we wish for ourselves. These arguments – though still tentative in their development – signal new and sanguine directions for meta-ethical thought.

6 Korsgaard's *The Sources of Normativity* (1996)

6.1 Introduction

Thus far, this dissertation has focused on developing a positive account of normativity. This Chapter marks a turning point, in that we move towards examining existing arguments in the literature. All of the arguments that I discuss in Chapters 6 through 10 profess to vindicate objective moral requirements through some kind of constitutivist, constructivist, or internalist basis. Since our project was motivated at the outset by the prospects of vindicating moral objectivity without appealing to theism or realism, it behooves us to examine each of these accounts, which are close neighbors to the project that we have pursued.

I begin in this Chapter with Christine Korsgaard's *The Sources of Normativity* (1996), one of the seminal works of recent moral philosophy. In *Sources*, Korsgaard argues for the following conclusion: if we value anything at all, then we must value the humanity of all persons.[156] Her general argument for this conclusion can be put as follows:

1. If an agent values anything at all, then she must have some *reason* for action, since she is conscious of the grounds of her choice. (*Reflective Premise*)

2. If an agent must have some *reason* for action, then she must have some practical identity or set of practical identities. (*Identities Premise*)

3. If an agent must have some practical identity or set of practical identities, then she must value her own *humanity* – her identity as being able to choose or value at all. (*Humanity Premise*)

[156] As Korsgaard puts it, "if we take anything to have value, then we must acknowledge that we have moral obligations" (1996: 92).

4. If an agent must value her own humanity, then she must value the humanity of all persons. (*Generalization Premise*)

5. Conclusion: If an agent values anything at all, then she must value the humanity of all persons.

This argument, with its various steps, has been examined at some length in the recent literature.[157] My aim in this Chapter will therefore be to surgically focus on one step in particular – the *Humanity Premise* – so as to raise some fresh points about it. I shall argue that this premise relies upon a *non-sequitur*.

To make that objection, I first begin with an explication of Korsgaard's argument for the *Humanity Premise* (§6.2). I then examine two recent objections to Korsgaard's argument for this step, from Street (2012) and Bukoski (2018), and I suggest that Korsgaard can find the resources to successfully respond to these objections (§6.3). Finally, I formulate the objection that Korsgaard's argument rests upon a *non-sequitur* (§6.4). To close, I suggest that because the *Humanity Premise* is pivotal to the success of Korsgaard's overall argument in *Sources*, our reasons to doubt this premise also constitute reasons to doubt the success of the overall argument.

6.2 Valuing Humanity: Korsgaard's Argument

The *Humanity Premise* is an intermediate step needed before Korsgaard can arrive at her intended moral conclusion. In what follows, I focus solely on Korsgaard's argument for the *Humanity Premise*, therefore setting aside questions about the subsequent *Generalization Premise*.[158]

Why does Korsgaard think that we must value our *own* humanity? Korsgaard begins with the idea that we cannot help but act for *reasons*, and this

[157] For discussion of Korsgaard's argument, see Cohon (2000), Kerstein (2001), Gert (2002), Regan (2002), FitzPatrick (2005), Langton (2007), Stern (2011), Street (2012), Theunissen (2018), and Bukoski (2018).
[158] We turn to questions related to the *Generalization Premise* in Chapter 7.

is because we have a reflective, self-conscious nature – one in which we are aware of the grounds upon which we act:

> The reflective mind cannot settle for perception and desire, not just as such. It needs a *reason*. Otherwise, at least as long as it reflects, it cannot commit itself or go forward. (Korsgaard 1996: 93)

This is what I have called the *Reflective Premise*, and it seems best understood as a claim about what it is like to be human: self-conscious choice is our *plight*.[159]

Next, Korsgaard argues that, when we act for reasons, we do so under the guise of broader practical 'identities':

> The reflective structure of the mind is a source of 'self-consciousness' because it forces us to have a *conception* of ourselves… a description under which you find your life to be worth living and your actions to be worth undertaking. So I will call this a conception of your practical identity… and all of these identities give rise to reasons and obligations. (Korsgaard 1996: 100-101)

This is what I have called the *Identities Premise*, and it seems best understood as a conceptual point about what it is like to act as a human: we understand our particular reasons in light of our more general identities.[160] In other words, we take ourselves to have reasons, in virtue of our being students, teachers, parents, citizens, or otherwise.

Then, Korsgaard points out that, because these identities are ultimately chosen by us – normativity is *mind-given* for Korsgaard – we could always discard any one of them:

---

[159] Later, in *Self-Constitution*, Korsgaard clarifies that this point is not concerned with logical, causal, or rational necessity: it is simply "our *plight*: the simple inexorable fact of the human condition" (2009: 1-2).

[160] Strictly speaking, I do not think this step is necessary for Korsgaard's argument to move forward, as she could just make the point that we need *reasons* for action, and that point would be enough to bring us to the subsequent stage of the argument.

> Because these conceptions are contingent, one or another of them may be shed. You may cease to think of yourself as a mother or a citizen or a Quaker. (Korsgaard 1996: 120)

This leaves us with a question: where are we to find an ultimate source of reasons? For we could always discard any particular identity as a source of reasons; but because we are reason-needing creatures (*Reflective Premise*), we will always be in a state of needing some source of reasons. Here is how Korsgaard makes this point:

> What is not contingent is that you must be governed by *some* conception of your practical identity. For unless you are committed to some conception of your practical identity, you will lose your grip on yourself as having any reason to do one thing rather than another – and with it, your grip on yourself as having any reason to live and act all. (Korsgaard 1996: 121)

This leads Korsgaard to the thought that we need some reason to choose *any* identity at all – some most *general* reason to adopt reasons in the first place. This reason, Korsgaard says, cannot come from the value we place on some *particular* practical identity, but only from the value that we place on our own *humanity*:

> But *this* reason for conforming to your particular practical identities is not a reason that *springs from* one of those particular practical identities. It is a reason that springs from your humanity itself, from your identity simply as *a human being*, a reflective animal who needs reasons to act and to live. And so it is a reason you have *only* if you treat your humanity as a practical, normative, form of identity, that is, *if you value yourself as a human being*. (Korsgaard 1996: 121, emphasis edited)

This is how Korsgaard reaches the *Humanity Premise*, which she more concisely puts as follows: "since you cannot act without reasons and your humanity is the source of your reasons, you must value your own humanity if you are to act at

all" (1996: 123). To illustrate Korsgaard's point, we could therefore imagine a conversation like the following. 'Why are you a Quaker?' 'Because I have chosen to be a Quaker – I have placed a value on it.' 'But couldn't you always discard being a Quaker? Don't you have an ultimate source of reasons, which cannot be so easily shed?' 'Good point. It is because I place a value on my very *capacity* for choice – my *humanity* – that I have reasons to choose any identity at all, and therefore I have an ultimate source of reasons.' There is hence a 'regress' of normative questions that could be raised, as to what ultimately gives us any reasons at all, and Korsgaard believes that *because we need reasons,* even if we shed any particular identity, this forces us to value our own *humanity*.

Let me now turn to examine to recent objections to this argument, from Street (2012) and Bukoski (2018), which will also help us to elucidate the argument further.

6.3 Valuing Humanity: Street and Bukoski

We begin with Street's objection. In order to understand it, let us suppose that one is the sort of person who is questioning one's own reasons to maintain some practical identity. For instance, suppose that one has decided to abandon one's practical identity as a Quaker. Now, one needs some new practical identity to serve as a source of reasons, though one could abandon that identity as well. Eventually, one must ask the following question in the normative 'regress':

> Does it really matter whether we act as our humanity requires, whether we find some ways of identifying ourselves and stand by them? (Korsgaard 1996: 123)

Korsgaard's response is that we *must* answer yes to this question, because we are reason-needing creatures – we cannot abandon the fact that we need reasons.[161] That leads Korsgaard to conclude that we must value our own

[161] I am setting aside here Korsgaard's discussion of suicide and complete normative skepticism (1996: 160-164).

humanity: our identity as "a reflective animal who needs reasons to act and to live" (1996: 121).

Street's objection is that there is a certain point in this 'regress' where the question that is raised becomes *illegitimate*. This is because, as Street suggests, there is no perspective from which to ask and answer this ultimate question, *other than from the perspective of some particular practical identity or set of reasons that we have already adopted.* So, to illustrate this point, Street envisages that we are asking ourselves this ultimate question in the regress: "Is there any reason why I should take something or other to be a reason?" (2012: 49). Street writes that we should "reject the question at hand as ill-formulated" (2012: 49). This is because, since Korsgaard has rejected mind-independent normative standards, we can only ask and answer normative questions *from within particular substantive standpoints that have already been adopted*. Here is how Street puts this point:

> According to metaethical constructivism, however, the standards that determine whether something "really matters" are set by one's own judgments about what matters; the standards that determine one's reasons are ultimately set by one's own judgments about what count as reasons; and there *are* no standards apart from this: this is the rejection of realism. It is therefore *illegitimate* to stand apart from every last one of one's judgments about what matters, and then to ask whether something further matters; it is *illegitimate* to stand apart from all of one's judgments about what counts as reasons, and then to ask whether one has some further reason. (Street 2012: 49, emphasis added)

Street's point is that normative questions about our reasons only make sense from within some practical standpoint in which we have *already* adopted certain substantive values. One could reject being a Quaker, but then one would need to adopt some new practical identity, in order for questions about reasons to make sense – there is no standpoint that abstracts away from *all* identities, from which one could then ask what reasons one has to adopt some identity. For Street, then, the regress of normative questions "comes to an end not with any substantive

value, but with an understanding of the exact moment at which normative questions cease to make sense – namely, the moment one divorces oneself from the practical point of view altogether, refusing, for that moment, to take any value for granted" (2012: 51-52).

It seems to me, however, that it *is* possible to imagine a standpoint from which we can abstract from all particular substantive values – or at least to suspend our commitment to them, *as reason-giving for us* – yet still ask questions about our normative reasons; and this would be the rejoinder that I imagine we could take on behalf of Korsgaard. To give this proposed standpoint a name, let me call it the standpoint of *freedom* – the internal standpoint from which we suspend or abandon particular substantive commitments as reason-giving, yet still retain an ability to assess different normative possibilities and make choices. Street claims that questions about reasons are only legitimate from within a standpoint in which we have already adopted certain substantive values – for instance, the thought would be that I can only assess how great being a Quaker is, from the perspective of *being* a Quaker, or from some other chosen identity. But it seems to me that these are not the only available options; for humans are capable of *suspending* their particular substantive commitments – 'I am not committed to that value, as a source of reasons for me' – yet still imagine, entertain, and ultimately *choose* some course of action. In this envisaged standpoint of freedom, we might speculate about what various normative possibilities are like – we might even imagine ourselves wearing the 'hat' of one – but this does not mean that we have already *committed* to the reason-giving status of any option.

Insofar as such a standpoint of freedom is cogent, we can see how Korsgaard could respond to Street. As both parties agree, we can question and then discard particular substantive commitments as reason-giving. But what happens when we question whether we have a reason to make a choice at all, from within the standpoint of freedom? Unlike with particular substantive commitments, Korsgaard can insist that we cannot answer *this* question except in the affirmative: for unlike particular practical identities, which we can

abandon the authority of, we cannot abandon the authority of the standpoint of freedom, since we *must* make a choice. And that is because we cannot avoid being *agents* – so long as remain alive as humans, we cannot help but act for *some* reason, and therefore to confer an authority on our general capacity for choice.[162]

If that is right – if it is conceivable that humans can step back from their particular substantive commitments to a suspended standpoint of freedom – then we can respond, on behalf of Korsgaard, that it is *not* illegitimate to keep asking normative questions in the style of a regress. For the question of whether we have a reason to choose any identity at all, after suspending all particular substantive commitments, *is* one that we can ask, from within the standpoint of freedom. Korsgaard's point is simply that we will inevitably find ourselves affirming the normative authority of our capacity for choice – and in her account, thereby valuing humanity.[163]

Let me move to consider a second objection from Michael Bukoski (2018), who argues that Korsgaard's account is ultimately *circular*. According to Bukoski, Korsgaard does not explain why valuing humanity is what can *exclusively* end the normative regress, rather than valuing other practical identities:

> Korsgaard needs to show that valuing humanity can end the reflective regress in a way that other potential justifications cannot. It would be circular to assume in the course of doing so that the value we place on other practical identities derives from the value we place (or are committed to placing) on humanity. If that claim is true, then it is true *because* valuing humanity can end the reflective regress, so it cannot without circularity be used in an argument to *show* that valuing humanity can end the reflective regress. (Bukoski 2018: 221)

---

[162] Again, this is setting aside questions about suicide or complete normative skepticism (see Korsgaard 1996: 160-164).

[163] It is important for me to emphasize that the suggestion I have made here about 'the standpoint of freedom' is still a *practical* standpoint, in the sense that I discussed earlier (§1.3): it is a *phenomenally conscious* standpoint from which we are sensitive to normative considerations, and can think imagine and think through them. Importantly, the standpoint is something entirely *mind-dependent*.

Bukoski's point is that Korsgaard has not shown why valuing humanity can *exclusively* end the regress, rather than "other potential justifications." Bukoski then suggests that valuing humanity fares no better than valuing other practical identities in ending the regress:

> Korsgaard has not shown that valuing humanity can end the reflective regress. Like any other practical identity, we can question whether it is justified and reject it if no adequate justification is forthcoming. A reflecting agent, if she is to act, must either find some other solution to the reflective regress, or if there is none, endorse some practical identity or other as sufficiently well-justified despite the possibility and coherence of further questioning. (Bukoski 2018: 221)

Bukoski implies that valuing humanity is just like valuing any other practical identity: we can "question whether it is justified and reject it if no adequate justification is forthcoming." So, Bukoski claims that Korsgaard has begged the question – she has assumed the very conclusion that she needs.[164]

Let me now suggest that Korsgaard has the resources to respond to this objection: for if we look at the text of *Sources*, we can reconstruct on Korsgaard's behalf an *argument* as to why humanity is the *exclusively* authoritative end to the normative regress, unlike other practical identities. Again, this is a point that we have already seen: because the authority of other practical identities can be abandoned, but the authority of *choice* cannot be abandoned, it has a distinctive power in ending the regress. Here is Korsgaard on why we cannot abandon the normative authority of our own *humanity*:

> You must value your own humanity if you are to value anything at all. Why? Because now that you see that your need to have a normative conception of yourself comes from your human identity, you can query the importance of that

[164] See Street's similar comments about Korsgaard's "question-begging reply," though Street emphasizes that this is not what she sees as the "real problem" (2012: 48-49).

> identity. Your humanity requires you to conform to some of your practical identities, and you can question this requirement as you do any other. Does it really matter whether we act as our humanity requires, whether we find some ways of identifying ourselves and stand by them? *But in this case you have no option but to say yes.* Since you are human you *must* take something to be normative, that is, some conception of practical identity must be for you. If you had no normative conception of your identity, you could have no reasons for action, and because your consciousness is reflective, you could then not act at all. Since you cannot act without reasons and your humanity is the source of your reasons, you must value your own humanity if you are to act at all. (Korsgaard 1996: 123, emphasis added)

On the basis of this passage and others, we could reconstruct an argument on behalf of Korsgaard as to why humanity is the *exclusively* authoritative way to end the regress:

1. Humans need a way of ending the regress of normative questions: we must find an ultimate source of reasons.

2. It is not possible to stop the regress by seeing *particular* practical identities as ultimate sources of reasons, because it is always possible to question and abandon these particular identities.

3. It *is* possible to stop the regress with *humanity* as an ultimate, authoritative source of reasons, because a human cannot legitimately abandon *humanity* as a source of reasons – that would be to abandon being a self-conscious agent at all.

4. There is no other way to end the regress.

5. Conclusion: Therefore, humans must regard their own humanity as an exclusively authoritative, regress-stopping source of reasons.

This argument, insofar as it is what Korsgaard has in mind, is not a *circular* one. For it never assumes its conclusion as a premise: the argument never just takes for granted the idea that humanity is the exclusively authoritative end to the regress, but rather tries to bring us to that conclusion. Korsgaard does explain why the normative authority of humanity specifically resists abandonment, unlike the normative authority of particular identities: it is because abandoning the former would violate our ability to have any reason at all, unlike abandoning particular identities. In our example, you can stop being a Quaker without losing all of your reasons for action; but if you stop regarding your capacity for choice as authoritative, then you will lose all reasons for action. This is what makes *humanity* – our capacity for choice – constitutively harder to abandon as an authoritative identity, for an agent who needs reasons. So, when Bukoski writes that "Korsgaard needs to show that valuing humanity can end the reflective regress in a way that other potential justifications cannot," we can imagine that Korsgaard can respond that she *has* done this: she has argued that only humanity can end the regress, unlike particular practical identities, because it is distinctively hard to abandon.

I have just suggested that Korsgaard's account need not involve illegitimate questioning, nor does she offer an argument that is circular. But let me now suggest that there remains a problem with her argument: it relies upon a *non-sequitur*.

6.4 The Problem: A Non-Sequitur

Korsgaard may claim that we can question particular substantive commitments from a more general standpoint of freedom, and she is right to claim that it is harder to constitutively abandon the authority of that more general standpoint – from which we make free *choices* – than it is to abandon the authority of particular identities. But now, let me suggest that the *Humanity Premise* nevertheless seems to rest upon a *non-sequitur*. For, from the fact that we must see our capacity for choice as an *authoritative*, ultimate source of reasons – even

an *exclusively* authoritative one – does not entail that we must also *value* that capacity, in a substantive sense.

Recall the ultimate point of Korsgaard's argument: it is to bring an agent into *moral* requirements. This means that valuing humanity must take on some kind of substantive character: for instance, at the end of the argument, we are meant to have substantive obligations to respect the humanity of all other persons, i.e. their capacity to choose. This means that the *Humanity Premise* must be understood as seeking to demonstrate that we have substantive obligations to respect – say, to protect or not interfere with – our own capacity for choice. But just because we see our humanity *as an exclusively authoritative source of ultimate reasons* does not entail that we must *value it as a substantive end*, and this is where the non-sequitur lies: the latter simply does not follow from the former.

Consider an example. Suppose that you have decided to fight in the war, knowing that you might die. When you ask yourself what gives you reason to fight, you conclude that it is ultimately because you *chose* to fight – and hence you recognize the normative authority of your humanity, i.e. your capacity for choice. But this would not entail that you must therefore *value* your own humanity, in a substantive sense: it would not entail that you must seek its preservation or prevent it from being lost in the war.[165]

The philosophical point is a general one: to regard X as an authoritative source of reasons does not, on its own, entail any substantive obligations towards X.[166] Here is a second example. Suppose that, having enlisted in the war, you regard your commanding officer as an authoritative source of reasons for you: *whatever* he or she says, you will follow, as more authoritative than any other possible source of reasons. But on some fateful day in the war, for strategic or other reasons, the commanding officer asks you to shoot him or her. Insofar as

[165] Street also makes a point along these lines, when she suggests that a 'rational' worker insect – even if she acknowledges that her reasons stem from her own capacity for choice or valuing – is not committed to the conclusion that she cannot sacrifice her life for the Queen insect (2012: 52-55).

[166] See Langton: "In general we don't think the source of something valuable must itself be valuable" (2007: 176). Cf. Korsgaard (2021: 172).

you regard this as an ultimately authoritative command, then you would be obligated to shoot, *rather than to value the life of your commanding officer as a substantive end*. Here, the 'X' is the commanding officer, whereas the point at stake in Korsgaard is one in which 'X' is your own humanity: just because you must see your humanity as authoritative does not mean that you must take on substantive obligations towards it.

Korsgaard might respond that, because the continued existence of our humanity is what sustains the value of what we care about, our caring about that particular thing entails that we must protect the continued existence of our humanity – so as to 'keep the flame alive' in the source, as it were. This might be contingently true in particular cases, insofar as the accomplishment of some specific goal requires our staying alive and continuing to be able to choose. But it would still not entail any *necessary* obligation to stay alive and preserve our capacity for choice, as the cases of war illustrate. For again, even if our reasons for going to war are sustained by our capacity for choice, our reasons for going to war might *involve the need to accept death*. More generally, even if we had to give up our capacity for choice in order to save the life of someone we cared about, our reason to save the life of someone we cared about would come from our capacity to choose, but it would not entail that we must protect our capacity to choose.

Because Korsgaard's argument rests upon a non-sequitur, it does not establish that, insofar as we value at all, we must value our own humanity. We can see the non-sequitur in the following statement: "it is a reason you have only if you treat your humanity as a practical, normative, form of identity, *that is*, if you value yourself as a human being" (1996: 121, emphasis added). To treat one's humanity "as a practical, normative, form of identity" is to see it as an authoritative source of reasons for oneself; but that is not the same as *valuing* yourself as a human being, in a substantive sense.[167]

[167] Interestingly, note that, in her most recent work, Korsgaard seems to amend her claim to de-emphasize humanity as a substantive end: she writes that we do not need "the claim that the power of rational choice is a valuable property," but only the claim that "in virtue of the power

Conclusion

This Chapter has focused on Korsgaard's pivotal claim that we must value our *own* humanity. I have suggested that, although Korsgaard may offer rejoinders to two recent objections to this claim, she must still reckon with the objection that her argument rests upon a non-sequitur. Since her account of morality in *Sources* hinges upon this pivotal claim, we have therefore raised an important question for her overall meta-ethical position.[168]

of rational choice, we assign ourselves a normative *standing* – the standing to legislate the value of our own actions and ends" (2021: 184, emphasis added).

[168] Some of the arguments in Chapters 6 and 9 first appeared in my BPhil essay, "Valuing Humanity in Kantian Constructivism," though the arguments have been updated, expanded, and re-written here.

7 Korsgaard's *Self-Constitution* (2009)

7.1 Introduction

This Chapter turns to examine Korsgaard's later position in *Self-Constitution* (2009). This work aims to establish that "the kind of unity that is necessary for action cannot be achieved without a commitment to morality" (2009: xii). It can therefore be read as arguing that moral requirements are *constitutive* of human action.

In comparison to the argument from *Sources* (1996), this later position involves a stronger emphasis on constitutive explanation, as well as on the idea of being an internally *unified* agent.[169] Much of *Self-Constitution* is focused on establishing the idea that, because we must assume a unified agent behind our actions, we must only act upon *universalizable* laws, i.e. laws that conform with the Categorical Imperative. However, Korsgaard makes it clear that this universalizability requirement does not yet bring us to a moral conclusion, unless we can establish a further proposition as well: that we must treat reasons as public. For as Korsgaard, emphasizes, "it is *only* on the public conception of reasons that a universalizability requirement is going to get us into moral territory" (2009: 192, emphasis added).

This creates an elegant dialectical situation for our purposes, in that it allows us to focus on engaging one specific claim: that we must regard reasons as public.[170] I begin by examining Korsgaard's defense of this claim in *Self-Constitution*, arguing that her ultimate justification for the claim seems insufficient (§7.2). I then ask whether Korsgaard can appeal to an earlier argument from *Sources* to supplement this defense, and I conclude that her

[169] Korsgaard herself writes that *Self-Constitution* and *Sources* argue for "the same conclusion," but that *Self-Constitution* does so "with a more direct focus on agency – on what is necessary to constitute yourself as the author of your actions" (2009: 25).

[170] There has been less critical examination of *Self-Constitution* in the recent literature, in comparison to *Sources*. For instance, even influential articles written *after* Self-Constitution – e.g. Street (2012), Bukoski (2018), and Theunissen (2018) – primarily focus on the earlier argument from *Sources*.

earlier argument faces problems as well (§7.3). I conclude that, because Korsgaard's overall position in *Self-Constitution* hinges on the claim we must treat reasons as public – a claim which is not sufficiently defended – we have grounds to question the success of the overall argument.[171]

7.2 The Publicity of Reasons: Korsgaard's Argument from *Self-Constitution*

Let me begin by again emphasizing the dialectical place of the publicity claim – the claim that we must treat reasons as public – in the overall argument of *Self-Constitution*. Much of the book is focused on establishing the conclusion that human agents must accept some kind of universalizability requirement.[172] But Korsgaard emphasizes that the universalizability requirement does not yet establish *moral* requirements, because it is compatible with a *private* conception of reasons:

> As many philosophers have pointed out, the privacy of reasons is consistent with a kind of universalizability requirement. If I conceive of reasons as private, and accept a universalizability requirement, I am committed to the view that… if I think that the fact that something will make me happy is a good reason for *me* to do it, then universalizability requires me to think that the fact that something will make you happy is a good reason for *you* to do it. But my happiness is still mine, and yours is still yours; mine is a source of reasons for me, but not for you; yours is a source of reasons for you, but not for me. (Korsgaard 2009: 191)

If we treat reasons as private, then we are not yet forced to treat others' reasons *as our own reasons*, since "a private reason is a reason whose normative force is private, in the sense it belongs to only one person" (2009: 191). This is why

[171] Indeed, our arguments against publicity also constitute a further reason to doubt the earlier argument from *Sources* as well, targeting the *Generalization Premise* (see §6.1).
[172] Korsgaard does not turn to the publicity of reasons until the penultimate Chapter 9.

Korsgaard must establish that we must treat reasons as public, in order for her overall argument to work: for it is "only on the public conception of reasons that a universalizability requirement is going to get us into moral territory" (2009: 192). If we treat reasons as public, then we would be acknowledging "reasons whose normative force can extend across the boundaries between people" (2009: 191). This would provide the following result:

> On the public conception I must take your reasons for my own. So if I am to think I have a reason to shoot you, I must be able to *will* that you should shoot me. Since presumably I can't will that, I can't think I have a reason to shoot you. (Korsgaard 2009: 192)

If reasons are public, and we have a universalizability requirement, then *your* reasons also have normative force *for me*: so if *you* have a reason to escape death, then *I* also have a reason to help you escape death. That is what would bring us into *moral* territory.

Let me now examine what Korsgaard's argument is for the claim that we must treat reasons as public, which comes in some of the crucial final sections of the book.[173] Her argument for publicity can be distilled into two essential claims:

1. *If* we engage in personal interaction, *then* we must treat reasons as public. (*Conceptual Claim*)

2. Human agents *must* engage in personal interaction. (*Interaction Claim*)

3. Conclusion: human agents *must* treat reasons as public.

The structure of Korsgaard's argument therefore runs as follows: she first outlines a conception of personal interaction, arguing for the *conditional* claim that, if we engage in this form of personal interaction, then we must treat reasons

[173] Specifically, §9.4 to §9.4 of *Self-Constitution*.

as public (what I call the *Conceptual Claim*); and she then tries to show that we *must* engage in that specific conception of personal interaction (what I call the *Interaction Claim*). I shall eventually suggest that Korsgaard offers insufficient defense for the *Interaction Claim*.

We can start with the *Conceptual Claim*, which Korsgaard articulates as follows:

> *If* personal interaction is to be possible, we must reason together, and this means that I must treat your reasons, as I will put it, *as reasons*, that is, as considerations that have normative for *me* as well as you, and therefore as public reasons. (Korsgaard: 2009: 192, emphasis added)

What makes Korsgaard think that this conditional claim is true?[174] She appeals to the thought that, if we do not treat another person's reasons *as reasons*, but simply as things to be navigated around (like obstacles), then as a conceptual matter, we would not really be *interacting* with that person:

> Suppose you and I are related as student and teacher, and we are trying to schedule an appointment. "Stop by my office right after class," I say, thinking that that will be convenient for me, and hoping that it will also be convenient for you. It isn't, as it turns out. "I can't," you say, "I have another class right away." So I have to make another proposal. It's important to see why I have to do this: it's because having the meeting is something that we are going to do together… To perform a shared action, each of us has to adopt the other's reasons as her own, that is, as normative considerations with a bearing on her own case. (Korsgaard 2009: 192)

Again, the point here is a conditional one: Korsgaard has not yet shown that the teacher *must* care about the student's schedule, *period*; rather, she has outlined a conception of personal interaction, and suggested that *if* two people are to

[174] She also writes: "suppose I don't treat your reasons as reasons, as considerations with normative force for me… if that's how it is, *no personal interaction is going to be possible*" (2009: 193, emphasis added).

interact in that sense, then they must treat each other's reasons as public. Korsgaard notes that the alternative would mean "that our relationship will be a kind of war, or combat" (2009: 194). As she writes:

> What's the alternative? Suppose I don't treat your reasons as reasons, as considerations with normative force for me… I don't try to work out something that's good for both of us, because your reasons are *nothing* to me, and any objections that might be grounded in your reasons are nothing to me… Obviously, we can't relate at all on those terms. So if that's how it is, no personal interaction is going to be possible. (Korsgaard 2009: 193)

This would not be personal interaction "in the Kantian sense" (2009: 200). Though someone else might *call* this interaction, they would be appealing to some different and thinner conception of interaction, in which we navigate one another like obstacles.[175] So, by outlining a particular conception of personal interaction, Korsgaard arrives at the *Conceptual Claim*: if we are to engage in personal interaction in this sense, we must treat each other's reasons as public.

The obvious question that then emerges is: why should an agent be interested in engaging in this conception of personal interaction? Why not just forego such personal interaction, if one is Caligula, who doesn't care for such interaction? To complete her argument, Korsgaard must show that human agents *must* interact in this specific sense – the *Interaction Claim*.

Let me now suggest that Korsgaard's defense for this claim is insufficient. First, I must confess to the reader that, although I have a deep and sustained regard for Korsgaard's work, this final stretch of argument is one that has been difficult for me to interpret, and so I will rely upon a combination of presenting her unedited text, alongside some attempted interpretations.

Korsgaard begins by acknowledging the obvious question here, which is why someone should be interested in personal interaction in her specific sense. As she writes:

[175] See Korsgaard (2009: 194).

> For of course the private reasoner will say that he can choose whether to interact, in the Kantian sense, with other people or not. He can decide he will engage in shared deliberation, with certain other people, for various private reasons of his own. But until this decision is made he has no reason to accord normative force to the reasons of other people. So what are these private reasons for personal interaction going to look like? (Korsgaard 2009: 200)

Korsgaard then offers three brief suggestions as to what an answer to this question could look like, before rejecting each of these options:

> Perhaps in view of certain ends that he and another person share, he thinks that the other person would serve well as an ally: they both can improve their efficiency by pursuing these ends together. Perhaps he and another person have made an agreement to take each other's reasons into account, as a kind of truce or social contract. Or perhaps he simply cares about some other people – he loves or respects them – and taking their reasons to be normative for him is part of what it means to care in this way. (Korsgaard 2009: 2000)

Korsgaard rejects the first suggestion, because it "does not really give me a reason to accord your reasons any normative standing of their own, although it may give me a reason to pretend to, in order to keep you on my side" (2009: 201); she rejects the second suggestion, because it runs us into "Hobbesian problems" where one could "free ride on the agreement" (2009: 201); and she rejects the third suggestion, because "it is not, therefore, something chosen" (2009: 202). If these three suggestions cannot answer the question as to why we should interact in Korsgaard's sense, then what can?

Korsgaard's answer – which comes in a final stretch of argumentation, crucial to the success of her constitutivist project[176] – rests upon the bold idea that interaction with *others* is just like interaction with *ourselves*:

[176] See Korsgaard (2009: 202-205).

> But for now, leave all those arguments aside. Let's suppose that you can just decide to treat someone's reasons as reasons, with normative implications for you, but that you need not do that unless you choose to. Would that show that morality is optional, depending as it does on whether you have any private reasons that favor personal interaction? It is not that simple, for there is one person with whom even the most determined private reasoner must interact in the way that Kant's theory requires. And that is himself. (Korsgaard 2009: 202)

Korsgaard claims that "even the most determined private reasoner" must interact with *himself* in the way that Kant's theory requires. But what does Korsgaard have in mind here? She continues to make the claim that the requirements for unifying one's agency *internally* are the *same* as the requirements for unifying one's agency with *others*:

> Personal interaction, I have argued, is quite literally acting with others. But for a creature who must constitute her own identity, it is equally true that acting is quite literally interacting with yourself. *The requirements for unifying your agency internally are the same as the requirements for unifying your agency with that of others*. (Korsgaard 2009: 202, emphasis added)

As the reader might surmise, this is an extremely important claim for Korsgaard – and it is also one that does not seem *obviously* true. Here is Korsgaard's defense of this final claim:

> But action is simply interaction with the self. If this is so, then respect for the humanity in one's own person, and the consequent treatment of one's own reasons as considerations with public normative standing, are the conditions that make unified agency possible… But couldn't he still will it as a public law only for himself, binding together only the parts – parts of the soul, or time slices, or whatever they might be – that are parts of himself? Couldn't he, that is, decide to respect only his own humanity? This is an ill-formed question. What is your own, in the individual sense of your own, is not your humanity

> but what you make of it, your practical identity, and the existence of that *depends* on your respect for humanity in general. (Korsgaard 2009: 204)

This passage – which could be interpreted in different ways – again presents the idea that one can only achieve *intrapersonal* unity, if one also has respect for the humanity of *others*. Here are two possible interpretations as to why Korsgaard thinks this is so: (i) because it is an "ill-formed" question to ask how certain requirements in the case of *intrapersonal* unity also entail requirements to treat *others'* reasons as public; or (ii) because the existence of one's practical identity *depends* on one's "respect for humanity in general." Let me now explain why both of these suggestions seem insufficient.

First, it does not seem like an *ill-formed* question to ask why *intrapersonal* unity also requires that we treat *others'* reasons as normative for oneself. Indeed, this seems to be the kind of question that motivates the overall investigation that Korsgaard is after: one might have thought that it was core to the entire project to try to demonstrate *why* this is indeed the case, i.e. to show why requirements of internal consistency within oneself also entail certain moral requirements towards others,[177] rather than to claim – at the end of the argument – that it is "ill-formed" to ask this question. For instance, Korsgaard writes that the Russian nobleman must not be at war with himself (2009: 203), but even if we accepted that, it still seems like a natural and *well-formed* question to ask why the Russian nobleman must therefore treat others' reasons as public. Rather than being an "ill-formed" question, such a question seems to lie at the heart – and chasm – of certain constitutivist efforts to vindicate objective moral requirements.

Second, the idea that the existence or value or one's own practical identity *depends* upon one's respect for humanity in general also does not seem obviously true. For, absent further defense, it seems like an open conceptual possibility that someone could adopt a practical identity for herself – as a Mafia member, or just as a successful egoist – while not having respect for the

[177] Plato's *Republic* could be read as trying to demonstrate such a result.

humanity in others. She does respect her *own* ability to construct a practical identity – and perhaps for this reason, grants a normative authority to her own capacity for choice – but at least absent further argument, it remains a genuine question whether this entails anything about necessary respect towards others. Caligula, or the unfeeling egoist, would insist that intrapersonal unity does not entail anything about respect towards others; and unless there is some further argument against the coherence of Caligula's position, Korsgaard's assertion that "the existence of [your practical identity] *depends* on your respect for humanity in general" leaves it an open question as to *why* we should accept this claim as true.[178]

This section has interpreted Korsgaard's argument in *Self-Constitution* as resting upon the claim that we must treat reasons as public. Furthermore, I have suggested that *that* claim relies upon a final assertion that the requirements of *intrapersonal* unity are the same as the requirements for unity with others – which seems to me an underdefended claim. Insofar as this exegesis and interpretation is accurate, then Korsgaard seems not to have provided an adequate defense in *Self-Constitution* for the pivotal claim upon which her overall account hinges.

7.3 The Publicity for Reasons: Korsgaard's Argument from *Sources*

In certain stretches of the argument we just examined from *Self-Constitution*, Korsgaard references her earlier argument for the publicity of reasons from *The Sources of Normativity*.[179] This raises the natural question: can Korsgaard resort to that earlier argument in order to achieve her intended result in *Self-Constitution*?

Let me now explicate Korsgaard's argument for publicity from *Sources*, before suggesting that it also receives an insufficient defense. In *Sources*,

[178] One possibility is that Korsgaard is presupposing, in this statement, some earlier arguments from *Sources* – in which case, we should turn to the examination of §7.3.
[179] See Korsgaard (2009: 196).

Korsgaard claims that the publicity of reasons is rooted in our social nature: "what both enables and forces us to share our reasons is, in a deep sense, our *social nature*" (1996: 135). She suggests that reasons are "inherently shareable" (1996: 135), and writes that "to act on a reason is already, essentially, to act on a consideration whose normative force may be shared with others" (1996: 136). She then offers a specific argument that is meant to parallel Wittgenstein's argument against the possibility of having a private language.[180]

Korsgaard's argument for publicity turns on the idea that our own demands on others to treat us well rest upon the idea that we are just *someone*, rather than upon some other contingent claim about ourselves. To make this point, Korsgaard imagines a situation of one person tormenting another:

> Suppose that we are strangers and that you are tormenting me, and suppose that I call upon you to *stop.* I say: 'How would you like it if someone did that to you?' And now you cannot proceed as you did before. Oh, you can proceed all right, but not just as you did before. For I have obligated you to stop. (Korsgaard 1996: 142-143)

Why has an obligation arose? Here, Korsgaard appeals to the idea that the tormenter would realize that his *own* claim towards others – to not be tormented – must rest upon the idea that he is "just human, just *someone*":

> How does the obligation come about? Just the way Nagel says that it does. I invite you to consider how you would like it if someone did that to you. You realize that you would not merely dislike it, you would resent it. You would think that the other has a reason to stop, more that he has an obligation to stop. And that obligation would spring from your own objection to what he does to you. You make yourself an end for others; you make yourself a law to them. But if you are a law to others in so far as you are just human, just *someone*, then the humanity of others is also a law to you. *By making you think these thoughts*,

[180] Korsgaard clarifies in *Self-Constitution* that her argument for the publicity of reasons is meant to be "analogous" to Wittgenstein's argument for the publicity of meaning (2009: 196n12).

> I force you to acknowledge the value of *my* humanity, and I obligate you to act in a way that respects it. (Korsgaard 1996: 143)

The idea that Korsgaard is appealing to here has strong similarities to the one that we defended in our own arguments for morality (Chapter 5).[181] But the obvious challenge to such a claim – as we had to address – is why the tormenter cannot merely think that *he* deserves not to be tormented, because of some more contingent reason: because of his personal identity, status, or the like. I offered two arguments for holding that a human cannot be committed to such a contingent claim, but must instead maintain a much wider claim concerning valuing *qua* valuing (see §5.2 and §5.3). But what is Korsgaard's defense against such an obvious challenge? She writes:

> You can try, of course. She says, 'My career is just as important to me as yours is to you, you know. I have ambitions too.' He says, 'It isn't the same thing for a woman.' What isn't the same? Does 'career' mean something different to her? Does 'ambition'? How about 'important'? Or (let's get down to brass tacks) how about 'I'? (Korsgaard 1996: 144)

However, it is not clear from this passage – as well as the remaining discussion that Korsgaard offers here[182] – what her argument is against the obvious challenge. For instance, in the above passage, it seems equally possible that Korsgaard could be appealing to some substantive norm – that it is wrong to appeal to one's gender or personal identity as the basis for a claim against being tormented – or that Korsgaard is suggesting that, once the tormenter is presented with these points ("by making you think these thoughts"), he will relent on his contingent claim, perhaps as an empirical matter. Either way, however, it seems

[181] The idea we defended was that the human practical standpoint is committed to the idea that our valuing deserves love, not in virtue of some contingent feature of ourselves, but because of the fact of valuing, i.e. valuing *qua* valuing. One difference between our claim and Korsgaard's is that appealing to 'valuing *qua* valuing' is even wider than appealing to one's *humanity* or *being human*, among other general differences in argumentation.
[182] See Korsgaard (1996: 144-145).

that more could be said here, such that it remains a vivid challenge to Korsgaard's account why someone *must* appeal to the fact that they are 'just someone' when they make claims on others not to torment them. Although I have offered some arguments earlier in an attempt to plug a similar meta-ethical hole, I am not confident that the text of *Sources* successfully does so, at least in the passages that I have surveyed.

I therefore conclude that, although Korsgaard makes some very helpful suggestions in *Sources* concerning the publicity of reasons, it does not amount to a sufficient defense of the claim that we *must* treat reasons as public.[183]

Conclusion

The argument for objective moral requirements from *Self-Constitution* rests upon the claim that we must treat reasons as public – and I have argued that, in both *Self-Constitution* and *Sources*, this claim does not receive a sufficient defense. In *Self-Constitution*, it is ultimately obscure why one *must* be interested in engaging in Korsgaard's conception of personal interaction; and in *Sources*, it remains an open question why our claims on others *must* be grounded in our being 'just someone.' It is therefore seems fair to ask whether Korsgaard can offer some further defense for the publicity of reasons, as the overall argument appears to hinge upon this point.

[183] See also Gert (2002) for critical discussion of the publicity claim from *Sources*.

8 Smith's Constitutivism (2013)

8.1 Introduction

In the last decade, Michael Smith has offered one of the most important attempts at explaining moral requirements in terms of what is *constitutive* of being an agent. In his own words, he has sought to show that we can derive "a substantive account of normative reasons for action" from "abstract premises about the nature of action and agency" (2015: 188). This, he has elsewhere said, is the "holy grail of moral philosophy" (2013: 9).

Smith's argument has the following structure to it: first, he tries to show that having certain desires – such as the desire to not interfere with others' capacities as an agent – is *constitutive* of being an ideal agent; and then, he claims that actual agents have reasons that are derived from the desires of their idealized counterparts (which I call his Dispositional Claim). The structure of this argument is clear in what it aims to accomplish: if the *ideal* agent must have desires that resemble moral requirements, and an *actual* agent has reasons that are grounded in the desires of the ideal agent, then we can show that actual agents have reasons that resemble moral requirements.

This Chapter offers a constructive and critical examination of Smith's account, focusing attention on both of these structural elements in his argument. First, I question whether Smith's derivation of the desires of the ideal agent are successful on constitutivist grounds – and I raise some reasons to be skeptical. Second, I argue that Smith has not adequately justified the Dispositional Claim – the claim that actual agents have reasons that are grounded in the desires of the ideal agent – on *constitutivist* grounds. This, I suggest, creates room for us to read Smith as a realist: although Smith *presents* himself as a constitutivist, who stands in contrast to the realism of Parfit (see Smith 2017), Smith's account ultimately rests upon a claim about normativity that risks looking like a brute and mind-independent normative fact. Because Smith's view is one that is still

in development, these remarks can be taken as raising concerns for the eventual view that is to be taken.[184]

I proceed as follows: after offering an explication of Smith's overall argument (§8.2), I turn to critically examine his derivation of the desires of the ideal agent (§8.3), followed by an analysis of how he attempts to defend the Dispositional Claim (§8.4). I conclude that Smith must buttress his account with further argumentation, in order for the view to succeed without appeal to brute and mind-independent normative facts.

8.2 Smith's Argument

As I mentioned, there are two basic parts of Smith's account: his claim that we can derive a notion of the ideal agent, with certain desires, on *constitutivist* grounds; and his claim that an actual agent has reasons that derive from the desires of the ideal agent (his Dispositional Claim).

How does Smith arrive at the first claim about ideal agents? He begins with the conception of an *agent* as "the sort of thing whose function is to realize its desires" (2015: 189). Smith believes that the accomplishment of this function involves two more specific capacities: a capacity "to have knowledge of that part of the world that concerns the ways in which it could act so as to satisfy its desires" (2015: 189), and a capacity "to put their intrinsic desires together with their knowledge of means to ends" (2015; 189). We can call these capacities for *knowledge-acquisition* and *desire-realization*.

With this conception of the *agent* in place, Smith then moves to constitutively derive a conception of the *ideal* agent. To do this, he appeals to what Judith Jarvis Thomson calls a "goodness-fixing kind" (Thomson 2008), which roughly refers to the kind of thing that, given its nature, entails an internal standard of assessment for better or worse instances of that thing. In Smith's words, goodness-fixing kinds "entail an internal standard of assessment against

[184] Smith's view is developed in a series of papers (see Smith 2013, 2015, 2020, 2021). As I understand it, the project is now being turned into a book manuscript.

which we can measure instances of the kind so as to come up with an ordering of the instances from best to worst" (Smith 2020: 62; see also Smith 2013: 17). Smith provides the example of a toaster: because it is part of the *concept* of a toaster that it warms and browns bread, this entails that a *good* toaster is one that does not burn the toast, whereas a bad toaster is one that does burn the toast (Smith 2013: 17). Smith thinks that the kind *agent* is like this: because the agent's function is to fulfill its desires through its capacities of *knowledge-acquisition* and *desire-realization*, an *ideal* agent is one that exercises these capacities to a maximal extent: "a good agent is someone who has and exercises, to a high degree, the capacity to know the world in which he lives and to realize his final desires in it" (2013: 18). Although questions could be raised here, I will take these points for granted and focus critical attention on later portions of Smith's argument.

Smith thinks that this philosophical structure then provides him the resources to claim that all ideal agents must have certain desires in common. More specifically, these are *dominant* and *final* desires – *dominant* in that they override all weaker desires (and other ones that are not constitutive of the ideal agent), and *final* in that they are concerned with something for its own sake.[185] Smith claims that there are at least three such dominant, final desires that every ideal agent must have: (1) a desire to not interfere with one's own capacity for knowledge-acquisition; (2) a desire to not interfere (and promote) one's own capacities for knowledge-acquisition and desire-realization in the future; and most importantly, (3) a desire to not interfere with (and promote) these capacities in other agents. It is this last one that brings us into moral territory. The derivation of each of these desires involves a specific stretch of argument, which I will turn to examine in the following section (§8.3).

Smith then closes his argument with the claim that an actual agent has reasons that are derived from the desires of the ideal agent – this is what I have called his Dispositional Claim. In his words, "what agents have normative reason to do is fixed by the desires of their ideal counterparts" (Smith 2020: 66;

[185] See Smith (2013: 10 & 23).

see also 2013: 18 and 2015: 187, 196). This is the claim that a real person in the world *inherits* her overriding reasons from the dominant, final desires of her idealized counterpart – where, apart from the aforementioned dominant, final desires, the ideal agent otherwise just desires what the actual agent desires.[186]

This argument therefore promises to land us in moral territory: for if actual agents have reasons that derive from the desires of the ideal agent – which includes a dominant, final desire to not interfere with (and promote) the agential capacities of others – then we will have reasons with substantively moral content to them. For instance, Smith suggests that these will involve reasons to keep promises, to care for our children, to help and courage others, to refrain from lying and cheating, and so on (see Smith 2013: 27-28).[187] Given that actual agents have reasons to perform such actions, Smith concludes that he has given morality a "rock solid foundation" (2013: 26).

8.3 The Desires of the Ideal Agent

Let me now move to examine Smith's derivations of the desires of the ideal agent, raising some constructive and critical comments. We begin with the first dominant, final desire: to not interfere with one's own capacity for *knowledge-acquisition.*

Smith argues that the ideal agent must adopt this dominant, final desire, because it is what allows the agent to avoid potential conflicts *between* her two capacities for knowledge-acquisition and desire-realization. For instance, imagine a case like the following:

> **Ancestral Belief.** Kevin desires to believe that he is a descendant of Confucius, whether or not he is indeed a descendant of Confucius.

[186] As Smith puts it: "Beyond this, each agent's ideal counterpart would simply desire whatever her potentially non-ideal self desires in those circumstances" (2015: 192).
[187] See Bukoski (2016: 137-146) for critical discussion of this point.

This is a case of someone who desires to believe that P is true, whether or not P is in fact true. But this pits an agent's capacities of *knowledge-acquisition* and *desire-realization* into conflict: for the maximal exercise of the capacity for knowledge-acquisition would instruct us to get clear on the facts, while the maximal exercise of the capacity for desire-realization would instruct us to come to form this belief at all costs. So, Smith notes that it is possible that the exercise of both capacities "pulls in opposite directions" (2013: 21), and we therefore need a "coherence-inducing" desire to rule out the possibility of such conflict in the ideal agent (2013: 22). Here is what Smith proposes:

> Imagine that all ideal agents have a dominant final desire that they do not now interfere with their current exercise of their capacity to have knowledge of the world… This would remove all potential for conflict between the full and robust exercise of an agent's capacity to know the world in which she lives and the full and robust exercise of her capacity to realize her desires. (Smith 2013: 23)

This is the requirement that the ideal agent *prioritize* knowledge-acquisition in cases of conflict – for it is through that prioritization that we "remove all potential for conflict" between the two capacities. So, in the example of **Ancestral Belief**, the ideal version of Kevin would have a dominant, final desire to get to know the facts, rather than to satisfy his desire to believe he is a descendant of Confucius.

The comment I would like to raise with this argument is that it seems to privilege knowledge-acquisition over desire-realization in an *ad hoc* or covertly substantive manner, rather than through constitutivist means. Consider that Smith's starting conception of the agent – before applying the idea of a 'goodness-fixing kind' – was defined as "the sort of thing whose function is to realize its desires" (2015: 189); its two capacities for knowledge-acquisition and desire-realization were only elucidated with reference to this more basic function. But nothing in this conception implies that, in a case of conflict, it is the capacity for knowledge-acquisition that should *win out* over the capacity for

desire-realization; in fact, when Smith introduces the capacity for knowledge-acquisition, he describes it as a capacity "to have knowledge of that part of the world that concerns the ways in which it could act *so as to satisfy its desires*" (2015: 189, emphasis added). On this framing of the capacity for knowledge-acquisition – which seems natural, given that this capacity is *derivative* of the function to realize one's desires – the function of acquiring knowledge is subservient to the aim of realizing one's desires. The *maximal* exercise of the capacity for knowledge-acquisition, when understood under this constitutive aim, would therefore not have priority over the desire, but would only be used to the extent that it could serve to bring the desire to fruition. In the case of **Ancestral Belief**, then, the capacity for knowledge-acquisition should only be used to the extent that it can help Kevin come to believe that he is a descendant of Confucius, which might include refraining from investigating the facts, after a certain point. The capacity does *not* seem to find constitutive grounds to take on a life of its own, in a way that overrides the desire that Kevin has to believe.[188]

Smith is aware of this objection, and he responds as follows:

> [One] could try suggesting that an ideal agent doesn't have to fully and robustly possess and exercise the capacity to know the world in which he lives, only the capacity to realize his desires in it, but then how would he be imagining that the ideal agent manages to realize his desires? (Smith 2013: 22)

The answer to Smith's question is straightforward: the ideal agent can realize his desire – say to believe that he is a descendant of Confucius – simply by refraining from seeking out the facts on the matter, after a certain point. As Kierkegaard imagines, someone could realize a desire to believe in God, through abstaining from further objective inquiry on the matter.[189]

---

[188] Of course, one can always claim that we simply do have normative reasons to believe in what is true – but unless such a claim is given a constitutivist defense, it will seem to assume *realism* behind it.

[189] I take this to be one of Kierkegaard's main points in *Concluding Unscientific Postscript* (1846).

Next, the second dominant, final desire that Smith argues for is a desire to not interfere with (and promote) one's own capacities for knowledge-acquisition and desire-realization in the future. Here are two passages in which Smith makes this claim:

> Imagine an agent who is now totally indifferent to the fact that she will later lack the capacity to realize her desire, or perhaps one who positively desires that she lacks that capacity. These agents are also less than ideal in virtue of not being now such as to be ideal later. The upshot of these parallels, according to the Constitutivist, is that ideal agents must also have dominant final desires that they do not interfere with their current or later exercise of their capacity to realize their desires, and to do what they can to ensure that they have the capacity to realize their desires later. (Smith 2013: 24)

> Given that they may well exist over time, ideal agents must also be in the present such that they can possess and exercise their epistemic and desiderative capacities in the future. Ideal agents must therefore desire not just that they do not currently interfere with the present exercise of their epistemic and desiderative capacities, but that they do not currently interfere with their future exercise of their capacities. (Smith 2015: 191)

From my reading, it is not entirely clear what kind of argument Smith has in mind in these passages. It seems that Smith is appealing to a certain kind of intuition that we might have that caring about one's future capacities makes one more 'ideal' – and certainly many of us do have this intuition. Indeed, it is this intuition that motivates some of Parfit's arguments in *On What Matters*, such as his *Agony Argument* (2011a: 73-82). The problem at hand is that Smith must be careful to explain why we must desire to not interfere with (and promote) our future agential capacities in *constitutive* terms – constitutive terms that root the requirement to have this desire back in the nature of his conception of the agent as a "goodness-fixing kind." It seems to me that, unless more is said, the passages above do not accomplish such a task. For instance, consider again that

Smith's original conception of the agent was "the sort of thing whose function is to realize its desires" (2015: 189). This definition does not make reference to future desires – the natural and presumptive understanding of it is that the agent is defined as the type of thing that seeks to realize *the desires it actually has*, i.e. its present desires. And note that the question of whether non-realists can show that we have reason to care about the future is a significantly debated one in the recent literature.[190] I would therefore suggest that more argument is needed for Smith to claim this dominant, final desire on constitutivist terms.

There is a response to this concern that Kathryn Lindeman (2019) presents on behalf of Smith. Lindeman draws a contrast between two types of human desires: ones that are *atelic*, i.e. which cannot simply be 'fully finished' at some clear end point, and others that are *telic*, i.e. ones that can indeed find some end point without simply renewing themselves. For instance, a desire to remain in good graces with my companion might be an *atelic* desire, whereas a desire to finish a paper would be a *telic* desire. Lindeman then claims that humans primarily have intrinsic desires that are *atelic*, rather than *telic*:

> After all, our intrinsic desires are for things that we find fundamentally desirable: caring for our family and friends, being a good neighbor, acquiring knowledge, maintaining our health. We do not engage in realizing our intrinsic desire in order to complete realizing our desires; in realizing our intrinsic desires, we do not aim at exhausting them – crossing them off a list, as it were. Because of this, the realization of intrinsic desires is not something that comes closer and closer to completion the more we realize those desires. (Lindeman 2019: 87)

If this is true – and it is the function of agency to realize our (atelic) desires – then we can see that being an agent is about realizing desires that have a persistent nature to them, one that requires our continual attention even in the

[190] Following Parfit's *Agony Argument* (2011a: 73-82), such discussion can be found in Sobel (2011), Shemmer (2020), and Lin (2020), as well as in my *Prolegomena to an Internalist Meta-Ethics* (2021: §3.3).

future. This is what allows Lindeman to claim that being an ideal agent, construed as such, necessarily involves future-concerned aspect to it:

> From this we can conclude that *the function of agency is atelic*… Because agentive capacities making ongoing contributions to an atelic function, there will be no point after which they need not maintain their ability to contribute. (Lindeman 2019: 87)

This leads Lindeman to claim that Smith's claim that the ideal must have desires concerning the future "is true of agentive capacities" (2019: 87).

There is much that is insightful in Lindeman's remarks here, but the problem is that the point she makes is ultimately a *contingent* one – it is a contingent observation that human beings have more atelic intrinsic desires than telic intrinsic desires, or that human beings have *only* atelic intrinsic desires. There is nothing built into the concept of being an *agent* – as "the sort of thing whose function is to realize its desires" (Smith 2015: 189) – which commands that the agent only have *atelic* rather than *telic* desires. Moreover, it is possible to find actual human agents for whom their strongest desires are telic, rather than atelic – for instance, just imagine a suicide bomber, whose primary goal in life has become oriented around completing his mission. So, Lindeman's suggestion here is helpful insofar as Smith might want to significantly reduce the ambition of his account, into one that rests upon contingent observations about *certain* human agents, but not *all* agents *qua* agents.[191]

Let me now turn to examining the third and most important desire of the ideal agent: the desire to not interfere with and promote the agential capacities of knowledge-acquisition and desire-realization in *other* agents. This is the most important desire for Smith to derive, for it is the one that leads to moral requirements – without it, all Smith would have shown is that we must care about our own ability to acquire knowledge, and our own future capacities. It is only

[191] After all, recall that Smith wanted to provide a derivation of substantive requirements from "abstract premises about the nature of action and agency" (2015: 188).

through demonstrating concern about *others'* capacities that Smith can bring us into moral territory.

Smith's argument for this final desire is based on observing certain *symmetries* between concern for *our own future* agential capacities and the agential capacities *of others*. As he puts it, "the reason for the extension is to acknowledge the following remarkable symmetries" (2015: 191). Here are those symmetries:

> Just as an agent's future possession and exercise of his epistemic and desiderative capacities may be dependent on what he now does, so the possession and exercise of other agents' epistemic and desiderative capacities may be dependent on what he does now, and just as his current possession of his epistemic and desiderative capacities is dependent on what his past self did, so his current possession and exercise of his epistemic and desiderative capacities is dependent on what other agents did. The possession and exercise of every agent's desiderative and epistemic capacities is thus thoroughly dependent on the possession and exercise of every other agent's desiderative and epistemic capacities, just as thoroughly dependent as any particular agent's future self's possession and exercise of his desiderative and epistemic capacities is on that of his current self. (Smith 2015: 191)

This passage points out certain parallels between our relation to our own future capacities and our relation to others. Given these parallels, Smith writes that "in order to be fully consistent, treating like cases alike, ideal agents must therefore treat other agents on a par with the way in which they treat future stages of themselves" (2015: 192).

My concern with this argument is similar to the one I raised about the second dominant final desire of the ideal agent: that it is not sufficiently clear how this is meant to be a *constitutivist* argument that roots the requirement for adopting this desire in the "goodness-fixing" nature of an agent. For instance, Smith points out that there are many parallels between the cases of relating to our own future capacities and relating to other people – but it is not clear how

pointing out these parallels entails a *constitutive* point, where the desire to care about our own future capacities *entails* a desire to care about others' capacities. Smith appeals to the notion of "treating like cases alike," but there is room to wonder how an appeal to this general idea amounts to a result concerning what is *constitutive* of agency, or constitutive of a desire towards our own future capacities. Many things share similarities, but that does not entail that a commitment to A – which may be similar *in some respects* to B – entails a commitment to B.

More deeply, we can object that there *is* a relevant discrepancy between the cases of caring about our own future capacities and the case of caring about someone else: the difference of personal identity. This is an important – and from the perspective of longstanding ethical debate, *cavernous* – difference between the two cases: one's own future person is still the same person as oneself, whereas other people are not the same as oneself. This constitutes a *prima facie* reason to think that the two cases are *not* in fact alike, and therefore a *prima facie* reason to think that a desire to protect one's own future agential capacities does not, on its own, entail a further desire to protect others' agential capacities.[192]

Because Smith's view is still in development, the ball remains in his court to develop further explications as to how these three dominant, final desires of the ideal agent are to be derived on *constitutivist* grounds. In the meantime, however, I have suggested some reasons to be skeptical that these desires have been successfully derived on such grounds.

8.3 The Dispositional Claim

In this section, I turn next examine the Dispositional Claim. This is the claim that links actual agents to the desires of the ideal agent that we have just examined: the Dispositional Claim maintains, in Smith's words, that "what

[192] See also Bukoski (2016: 134-134) for a similar point, and Lindeman (2019: 88-90) for some further discussion.

agents have normative reason to do is fixed by the desires of their ideal counterparts" (Smith 2020: 66). Without the Dispositional Claim, Smith would not have shown why the desires of the ideal agent are *relevant* to the normative reasons that actual agents have.

The Dispositional Claim is a claim about the nature of the reasons that we actual agents have – it claims that these are fixed by the desires of our ideal counterparts, rather than by some other source, such as our actual desires in the present, or by certain desire-independent facts about value, or the like. It is therefore a claim about reasons with *substantive* implications, and it is also a claim about reasons that *excludes* other possible claims about our reasons as being true. It is therefore a claim that beckons for some kind of justification or defense, unless it is intended as a brute fact about normativity.[193]

The principal problem is that, throughout Smith's corpus in developing his constitutivist account, he provides scant discussion as to what would justify the truth of the Dispositional Claim. Consider two places where Smith discusses the Dispositional Claim:

> The Dispositional Theory tells us that an outcome is finally good, as indexed to an agent A, just in case A's idealized counterpart – that is, A himself in the nearest possible world in which he has and exercises a maximal capacity to know the world in which he lives and realize his intrinsic desires in it – finally desires that outcome. (Smith 2013: 18)

> And I say that an outcome is more desirable-relative-to-us than another just in case our ideal counterparts – that is, we ourselves in the nearest possible worlds in which we are fully informed and our desires are maximally coherent and unified – prefer that outcome. This is the dispositional theory of value. Facts about what we have reasons to do are in this way fixed by the desires of our

[193] One could debate whether the Dispositional Claim is better understood as a first-order or second-order claim about our reasons. Either way, the point here is that it is a claim about our reasons, which *excludes* other claims as being true, and which comes with *substantive* implications – and it therefore requires some kind of justification, unless it is meant to be a brute normative fact.

> idealized counterparts, as these fix the facts about what's desirable-relative-to-us. (Smith 2015: 187-188)

In these passages and the lines that surround them, Smith asserts the Dispositional Claim as true, without providing an explanation as to why we should accept the Dispositional Claim as true – especially when, as I have noted, the Dispositional Claim is a claim about our normative reasons that excludes alternative possible conceptions of our reasons that other theorists champion. For this reason, Bukoski is correct to observe that "in the paper presenting his constitutivist theory of reasons, Smith provides *no explicit defense* of this dispositional constitutivist claim" (Bukoski 2016: 122, emphasis added).

At times, Smith seems to suggest the Dispositional Claim is true as an *analytic* claim. For instance, Smith at one point writes:

> There is an *analytic* tie between facts about what is desirable-relative-to-an-agent and facts about the desires of ideal agents… [and] there is an *analytic* tie between facts about an agent's reasons for action, on the one hand, and facts about which options that agent has and what is desirable-relative-to-that-agent, on the other. (Smith 2015: 193, emphasis added)

Later again, when referencing the Dispositional Claim, Smith writes that "these connections are all analytic" (2015: 198). But the problem is that the Dispositional Claim, insofar as it is taken as a claim about our normative reasons, in the ordinary sense, does not seem true as an analytic matter. The concept of a normative reason means something like 'a consideration that counts in favor of' – but this leaves open whether these reasons are substantively derived from our present desires, from our future desires, from the desires of some idealized version of ourselves, from some desire-independent value, or otherwise. The *concept* of having a reason does not seem to entail that we must follow the desires of our idealized selves. To see this, just imagine someone who perfectly understands the concept of a normative reason, but denies that she has any reason to follow the desires of her ideal counterpart, according to Smith's conception.

This person *knows* that the ideal agent would desire to not interfere with others' agential capacities, but she takes herself to have a reason to interfere, because she wants to gain something – and this seems to involve no confusion in the *concept* of a reason.

Otherwise, if Smith really has in mind some concept of a 'reason' that is analytically connected to the idea that our reasons derive from the desires of our ideal counterparts, then he would face the objection that this notion of a 'reason' – which we might call a *shmeason* – is one that does not have obvious normative import for agents who do not care about it. Why should we be interested in following these 'shmeasons'? It would be natural and legitimate to ask why this notion of a 'reason' is *normative* for us – and then Smith would face the question of normative justification again.

I have argued that Smith does not provide an adequate or explicit constitutivist defense for the Dispositional Theory. Absent such a defense, the Dispositional Claim – which makes an exclusive claim about our reasons, with substantive implications – might start to look like a claim about a brute fact concerning normativity. In other words, absent such a defense, Smith would need to be committed to realism.

Is that a problem for Smith? Any theorist is entitled to make a claim about there being brute facts about normativity – indeed, this is what realists do. But if this were the tack that Smith took, he would need to do so explicitly, and therefore to assent to the philosophical problems that are associated with realism (§1.2). Furthermore, on the basis of Smith's corpus, it does not *seem* like this is what he would like to do – for instance, he makes it quite clear that he is opposed to the kind of realism that Parfit and Scanlon defend, seeing "excellent reasons to reject Reasons Fundamentalism" (2017: 102), and he instead defends "Subjectivism" from the criticisms that Parfit makes against it (2017).[194] Moreover, Smith claims to proceed from some general Humean picture (Smith

[194] Adrian Moore, in interpreting Smith's earlier *The Moral Problem* (1994), associates him and Korsgaard as both opposed to realism: "neither Smith nor Korsgaard is trying to salvage the idea that our ethical views are views about what is out there independently of us" (2003: 5).

2021: 276), which is usually taken to be opposed to realism. And finally, given that Smith self-professedly claims to advance a *constitutivist* theory of normativity – one in which "normative facts of certain kinds are explained by facts about the constitutive feature of something" (2018: 371) – it would be natural to expect him to pursue an account that is constitutivist "all the way down." For if a constitutivist account is only constitutivist *in part*, mixing constitutivist claims with brute normative facts, then it would seem to ultimately sacrifice the standard philosophical reasons for pursuing constitutivism in the first place, which are to *avoid* the objections that beset claims about brute normative facts. Hence, Enoch insightfully writes that constitutivism is "the hope of generating a *comprehensive* account of (at least practical) normativity in terms of what's constitutive of action and agency. Constitutivists cannot help themselves to any reasons that are not accounted for by the constitutivist story itself" (Enoch 2021: 339). Insofar as Smith is motivated by the usual kinds of objections to realism, it behooves him to pursue an account of constitutivism that is constitutivist all the way down.

Conclusion

This Chapter has argued that Smith's account, at least in its current formulation, faces several concerns. First, I argued that his derivations of the desires of the ideal agent do not always seem to proceed on clearly constitutivist grounds; and second, I have suggested that his Dispositional Claim does not currently have a compelling justification on constitutivist grounds. Smith's account therefore risks a reliance upon realism. This is not necessarily a problem for Smith: but he must either come clean, or develop further constitutivist arguments.

9 Markovits' *Moral Reason* (2014)

9.1 Introduction

In this Chapter, we examine Julia Markovits' *Moral Reason* (2014). Markovits develops a Kantian account of morality which holds that, if we rationally will anything at all, then we must will the value of *humanity*: "there is one substantive end that we must all will, on pain of irrationality, if we will anything at all: that of humanity" (2014: 145). Like Korsgaard, Markovits attributes the inspiration behind her argument to Kant, claiming that she wants to "fill in and revise the Kantian argument" (2014: 131), and like Korsgaard, Markovits takes Parfit's position as one of her foils, claiming that she wants to make only *procedural* claims about rationality.[195]

I argue in this Chapter that Markovits' conception of procedural rationality faces two concerns. First, at a pivotal point in her argument, Markovits relies upon a procedural requirement of 'understandability to others' that does not seem to count as a procedural requirement, by her own lights. Second, because Markovits does not justify the *normativity* of her procedural requirements in constitutivist terms, but instead holds that they are "irreducibly normative" (2014: 55), her account seems to face the same objections that realism does.

I begin by explicating Markovits' overall account, focusing on the step of her argument that seeks to establish that we must value our own humanity (§9.2). I then raise the concern that the requirement of 'understandability to others' does not seem to be a procedural requirement, on Markovits' own terms (§9.3). Finally, I argue that Markovits appears committed to realism, and this raises some important theoretical questions about her project (§9.4).

9.2 Markovits' Argument

[195] See Markovits (2014: Chapters 1 to 3).

Here is how Markovits lays out her overall argument:

1. I value the ends I rationally set myself, and take myself to have reason to pursue them.

2. But I recognize that their value is only conditional: if I did not set them as my ends, I would have no reason to pursue them.

3. So I must see myself as having a worth-bestowing status.

4. So I must see myself as having an unconditional value – as being an end in myself and the condition of the value of my chosen ends – in virtue of my capacity to bestow worth on my ends by rationally choosing them.

5. I must similarly accord any other rational being the same conditional value I accord myself.

6. *So I should act in a manner that respects this unconditional value: I should use humanity (that is, rational nature), whether in my own person or in the person of any other, always at the same time as an end, never merely as means.*[196]

We will focus on the fourth premise, which I shall refer to as the *Humanity Premise*.[197] This premise maintains that, because one has a capacity to bestow value on one's own ends by rationally choosing them, one must see *oneself* as having an unconditional value – which Markovits uses interchangeably with the proposition that one must value one's own *humanity*. The *Humanity Premise* is therefore a premise that underpins the success of Markovits' overall argument, making it an important subject of examination.

[196] This is taken directly from Markovits (2014: 130).
[197] We used a similar name for a comparable premise in Korsgaard (see §6.2).

How does Markovits propose to establish the *Humanity Premise*? Markovits suggests that we reach it through certain *procedural* requirements of rationality, which she contrasts to *substantive* requirements:

> Internalism embraces a *procedural* conception of practical rationality, according to which the rational requirement to hold certain ends is generated indirectly by the relation of those ends to other ends we already hold, as a result, in particular, of requirements of internal consistency and coherence… According to an externalist, *substantive* notion of rationality, reason may require us to hold some (moral and prudential) ends directly, and regardless of what else is true about us. (Markovits 2014: 51-52)

Notice the following important lines from this conception: a procedural requirement is generated "indirectly by the relation of those ends to other ends we already hold," and they concern "internal consistency and coherence." The internalist, on Markovits' account, does not *directly* tell the agent what ends to hold, but rather is concerned with certain internal relations between her ends, given that she has already adopted for herself certain ends.[198]

Markovits then proposes two procedural requirements that are said to lead to the *Humanity Premise*: (1) systematic unity within our ends, and (2) understandability to others. The first requirement requires that our various ends must be internally coherent, in light of one another:

> There's just something arbitrary and dogmatic about valuing many such unrelated, unsystematic, contingently-chosen ends, without some more fundamental explanation for why they matter. (Markovits 2014: 132)

> The virtue of willing a mutually supportive, systematically justified set of ends is a virtue of procedural practical rationality… there is rational pressure on us,

[198] What motivates such an internalism? Markovits mentions epistemic humility and an analogy to theoretical reasons (2014: 54-64); in later work, she also discusses metaphysical, non-alienating, and practical considerations (Markovits & Walden 2021: 318-320).

> as Kant thought, to search for "an unconditioned condition" of value. (Markovits 2014: 133)

Markovits is claiming that it would be procedurally *irrational* to value different ends, without some ultimate end that gives value to all of them: for instance, it would be irrational to value brushing teeth, skipping in the sun, and reading philosophy, without some *fundamental* end that explains the value of these ends. Markovits therefore requires that we find an "unconditioned condition" of value – a point that I shall not dispute here, at least for the sake of argument.

Markovits then appeals to the second requirement of procedural rationality – that our ultimate end must be understandable to others – to arrive at the idea that we must adopt our own humanity as our ultimate end. Here is how Markovits presents the procedural requirement:

> So it's important that the end I recognize as the source of value, and so of systematic justification, for my ends *make sense* as a potential source of value for the ends of others. (Markovits 2014: 136, emphasis added)

Markovits also says that this ultimate end cannot be recognizable "only from a fully parochial perspective" (2014: 135). This is why, according to Markovits, *stamp-collecting* cannot legitimately serve as an ultimate end that we organize our other ends around:

> So much, then, for stamp collecting. It doesn't even provide systematic justification to *my* ends, *much less make sense from the perspective of anyone else's*… Even if stamp collecting became *all* I cared about, so that my own value-commitments looked quite systematically justifiable, I would fail terribly at demonstrating epistemically healthy humility. I would totally dismiss most other people's perceptions of value from the start, with no way of defending the dismissal. (Markovits 2014: 136, emphasis added)

Markovits claims that our ultimate end must be one that *makes sense* from the perspective of others' ends – and she therefore rules out stamp-collecting as an ultimate end that could be procedurally rational.

Markovits then mounts an argument via elimination, ruling out all possible ultimate ends as meeting these requirements, *except* for the ultimate end of valuing one's own humanity. So, for instance, Markovits argues that neither *happiness* nor *desire-satisfaction* can serve as our ultimate end (see 2014: 136-141); and having exhausted what she sees as the obvious options, she concludes that only *our own humanity* can serve as an ultimate end that both (i) systematically unifies our other ends and (ii) is understandable to others.

So, to show that we must value our own humanity – that we must take ourselves to be unconditionally valuable – Markovits proposes the procedural requirements of systematic unity and understandability to others, which leads her to the conclusion that we must adopt our own humanity as our ultimate end.[199] This is how Markovits proposes to arrive at the *Humanity Premise*, which then supports the remainder of her argument. Let me now raise some questions about the procedural requirement of 'understandability to others.'

9.3 Understandability to Others

Here is the first concern with Markovits' account: it is not clear that the requirement of understandability to others is a *procedural* requirement, on

[199] Interestingly, notice that Markovits' argument does not face the same *non-sequitur* that we said Korsgaard's argument faces (cf. §6.4):

> Kant's idea is *not* that, *because* we're the source of value of our chosen ends, we must *therefore* be valuable in ourselves. Rather, it's the *way* we bestow value on our ends that matters: we do this by being the more fundamental ends for whose sake we pursue our contingent ends. (Markovits 2014: 142, emphasis added).

Korsgaard essentially claimed the opposite, which is what incurred the non-sequitur:

> Kant saw that we take things to be important because they are important to us – and he concluded that we must *therefore* take ourselves to be important. (Korsgaard 1996: 122, emphasis added).

Markovits' own terms. Indeed, as I shall suggest, this requirement seems to implicitly appeal to substantive norms.

Recall that Markovits wants to rule out stamp-collecting as an ultimate end that is procedurally rational, because it does not "make sense" from the perspective of other's ends. Here, the point is not that it is conceptually impossible to imagine such a person; the point is rather that "even *if* stamp collecting became *all* I cared about… I would fail terribly at demonstrating epistemically healthy humility" (2014: 136). But here is where the concern arises. First, there is the straightforward question of why stamp-collecting is something that cannot "make sense" from the perspective of others' ends. But more deeply, there is a question as to why this requirement of understandability to others' is a *procedural* requirement of rationality at all. Recall Markovits' conception of a procedural requirement:

> Internalism embraces a *procedural* conception of practical rationality, according to which the rational requirement to hold certain ends is generated indirectly by the relation of those ends to other ends we already hold, as a result, in particular, of requirements of internal consistency and coherence… According to an externalist, *substantive* notion of rationality, reason may require us to hold some (moral and prudential) ends directly, and regardless of what else is true about us. (Markovits 2014: 51-52)

But if we reflect upon what understandability to others involves, it becomes difficult to see how this is a requirement that concerns an *internal* relation among our ends: one that is "generated indirectly by the relation of those ends to other ends we already hold." For the requirement of understandability to others establishes a metric of assessment that does not depend, in any way, on what other ends we ourselves hold – rather, it *outsources* the assessment of our ultimate end to a determination by others, as to what they find understandable or not. This creates a standard of rationality that has much less to do with an internal relation among our own ends, and one that has much more to do with an imposition from the outside.

If we pursue the point further, we might suggest that the *only* way to really determine whether some ultimate end is 'understandable to others' is to implicitly appeal to some *substantive* norm about what is valuable. Imagine the idiosyncratic character who organizes her life around stamp-collecting – all her time spent eating, sleeping, and interacting with others is done in the *service* of her greatest passion, i.e. cultivating the world's most beautiful stamp collection. Let us now ask the question: what could make it such that this ultimate end 'makes sense' to others or not?

The idea here does not seem to be that some *majority* of other people must actually approve of stamp-collecting as an ultimate end – this would be too contingent a claim to vindicate a categorical view about moral requirements, for a majority's views could always change.[200] But if understandability to others does not reduce to some kind of head count, then it becomes an open question what kind of metric we could appeal to, so as to determine what is understandable to others or not. Unless more is said, it seems reasonable to worry that, after all, the requirement of understandability to others here is making an implicit appeal to a *substantive* norm, concerning the first-order content of the ultimate end.

After all, we should ask: what *is* wrong with having stamp collecting as one's ultimate end? No doubt this would be a rare kind of character, and no doubt most of us do not have stamp collecting as our ultimate end. But we can imagine someone who is building her beautiful and impressive stamp collection, and who does not ever harm anyone else – she is taking great delight from it, and otherwise minding her own business. Although stamp-collecting is not *my* own ultimate end, I can profess that I find something endearing about this kind of character, rather than seeing a problem in her.[201] If we are to claim that this person is under a *normative* requirement to *change* her ultimate end to something different in *substance* – yet not explain this requirement of 'understandability to

[200] Indeed, Markovits writes that "the goal isn't, of course, to find an ultimate end that will accommodate everything individual people *happen to value*" (2014: 139, emphasis added).
[201] Unless this person is not pursuing her strongest *counterfactual* endorsements (see §3.4), or if this person is neglecting her own moral obligations towards others (see §5.4).

others' through claims about what others actually find valuable – then it seems we will be very close to making a directly substantive appeal to what *is* valuable, i.e. to do exactly what Markovits wanted to avoid, in setting aside externalism.

We have therefore raised a question as to whether Markovits' 'procedural' requirement of rationality is, after all, just as dogmatic as realism. One of the philosophical motivations that Markovits cites is that her account is less dogmatic than Parfit's. Indeed, with the stamp-collecting example, Markovits claims that it would be "very dogmatic" for someone to refuse to give up the ultimate end of stamp-collecting:

> But to insist that stamp collecting is an ultimate worth-bestower is *very* dogmatic... I would fail terribly at demonstrating epistemically healthy humility. I would totally dismiss most other people's perceptions of value from the start, with no way of defending the dismissal. (2014: 136)

But I must confess that I find the opposite to be the true: when the quiet stamp collector is pursuing her own beloved collection – harming no one else, and not forcing anyone else to collect stamps with her – it seems far *more* dogmatic to tell her that she must change her ultimate end, just because *others* are said to have trouble making sense of it. I conclude that Markovits' requirement of understandability to others, though presented as a procedural requirement, is one that comes far closer to appealing to substantive norms as to what is worth valuing as an ultimate end.

Let me also briefly note how Markovits' appeal to 'understandability to others' is different than how we employed the idea of intelligibility to others in our own **Conceptual Argument** for morality (§5.2). In that argument, we appealed to the idea that, because our valuing makes implicit claims on others to see our valuing as reasons for them, we constitutively seek to *reach* others with our claim – and that involves needing to formulate our claim in a way that is understandable or intelligible to them. This point that we made – which may seem to have a surface similarity to Markovits' account, in that both invoke the notion of understandability or intelligibility – is different for at least two

important reasons. First, there is no requirement in our argument that our ultimate *substantive end* must be one that others find intelligible; the idea is simply that, when we appeal to others to help us, or to see our valuing as having reason-giving status for them, we must *formulate* this claim towards them in a way that they can understand or find comprehensible. Indeed, *what* we are trying to make them understand is entirely different: it is a claim on them, rather than our own ultimate substantive end. Second, the argument I gave explicitly explains this requirement as one that is *constitutive* of our valuing, and it is therefore one that addresses and settles questions as to where the *normativity* of the requirement comes from. As we shall now see, that will prove to be an important difference from how Markovits conceives of the normativity of her requirements.

9.4 Lurking Realism

The second point about Markovits' account is that the normativity of her procedural requirements amounts to something like a brute and mind-independent fact – a point that Markovits herself concedes, but at rather subtle points in the book. This aspect of Markovits' view gives us grounds for suggesting that her account must face the same standard objections that realism faces.

To introduce this point, consider first the question that Crisp poses to all theorists who champion *procedural* requirements of rationality:

> If we allow that I should φ only in so far as a certain procedure has as its upshot that I should, do I have reason in the first place to act on the basis of the deliverances of the relevant procedure?... If the answer is affirmative, then we have an INE [an intrinsically normative entity] – a different INE from that referred to by the substantive realist, but an INE nevertheless. If the answer is negative, then we have nihilism. (Crisp 2006: 53)

Most mind-given theorists would attempt to answer this question in a way that does not appeal to the *bruteness* of normative requirements – i.e. to look for some *constitutivist* mode of explanation, as we have done. But perhaps surprisingly, Markovits' response to such a question is to dig in her heels, and to admit, explicitly, that the normativity of her procedural requirements *does* amount to something like a brute normative fact. For, in her own words, her standards of procedural rationality are "irreducibly normative" (2014: 55). Even more salient is a footnote in which Markovits explicitly contrasts her own account from Korsgaard's, emphasizing that the normativity of her procedural requirements does not come from *us*:

> We aren't, on the view I defend and the reading of Kant I am exploring, the *authors* of the law or source of its normativity, *because we are in no way the authors of the requirements of procedural rationality, or the sources of their normativity*. This is the sense in which the internalist view I defend, and (tentatively) attribute to Kant, is *non-reductive*… This makes my project less ambitious, in a way, than some Kantian constructivist projects, such as Korsgaard's: unlike Korsgaard, I am not trying to identify the "sources of normativity." (2014: 73n4, emphasis added)

As this passage makes clear, Markovits does not maintain that the normativity of her procedural requirements is mind-given, or constitutive of something we already do – indeed, in a more recent piece, Markovits distinguishes her own position from "more ambitious constructivists who want to give a constructivist analysis of normative facts quite generally" (Markovits & Walden 2021: 326).[202] What all of this makes quite clear is that Markovits is prepared to assent to the idea that there are brute and mind-independent normative facts – i.e. that there are normative *requirements* which are not constitutively grounded in our own practical standpoints.[203] Indeed, this appears to make Markovits a *realist* in the

[202] Markovits identifies Street as one such constructivist (Markovits & Walden 2021: 326).
[203] Indeed, note that Korsgaard herself conceives of procedural requirements quite differently from Markovits: for Korsgaard, procedural realism is simply defined as "the view that there are answers to moral questions" (1996: 35).

exact sense that she later defines the term: "For the realist, the 'hold' that normative facts have over us will be something of a brute fact" (Markovits & Walden 2021: 319).

Of course, Markovits is entitled to advance a position that is realist, as any other theorist is. The only problem is that she would then need to be prepared to defend her view from all the same kinds of objections that realism faces: these are the metaphysical, epistemological, dogmatic, and other objections that we mentioned earlier (§1.2). The tension in Markovits' account is that she both *professes* to want to avoid some of the objections that a theorist like Parfit would face, but still assents to the idea that there are irreducibly normative facts. For instance, in her recent piece, Markovits explicitly outlines some of the reasons that constructivism is more philosophically attractive than realism:

> First, [constructivism] seems to avoid some of the metaphysical problems that normative realism faces… Second, the dependence of the normative on our attitudes and activities that the constructivist proposes eases the threat of alienation. (Markovits & Walden 2021: 318)

But since Markovits assents to the idea that there *are* irreducibly normative requirements, it is an open and pressing question whether she can reap the theoretical benefits of a mind-given or antirealist view of normativity. For as I have emphasized earlier, the objections to realism take force once a theorist makes even *one* claim about brute normativity: for it is the fact that we are positing brute normative facts *at all* that raise metaphysical and other objections. It does not matter whether we call these requirements procedural or substantive – for if they are *normative* requirements upon us, then insofar as they are also posited as brute, they give rise to the standard objections towards realism.

Markovits seems aware that someone could ask why procedural rationality is preferable to substantive rationality, if both are irreducibly normative:

> Despite being irreducibly normative, the standards of procedural rationality may be both less controversial and, like their theoretical counterparts, significantly harder to question than the substantive standards of rationality to which externalists appeal – standards declaring us to be irrational simply in virtue of not caring about some end, like the well-being of others, regardless of that end's relation to our other ends and commitments. (2014: 55)

But as I have emphasized, when we think about the stamp-collector, it is not at all clear why the requirements in question are 'less controversial' than the ones that externalists or realists appeal to. The stamp-collector wants to retain her own ultimate end of stamp-collecting – that is what she loves – and none of her other existing ends are in contradiction with this ultimate end. Could it not be more 'alienating' to simply insist that she must change her ultimate end? Moreover, it is not obvious that such a standard of procedural rationality is "less controversial" than certain substantive requirements. For instance, is it less controversial to think that we must adopt an ultimate end that is not stamp-collecting, or to think that we have reasons to promote our own well-being, which is something that many substantive realists claim? Is it less controversial to think that stamp-collecting is not a legitimate ultimate end, or to think that we have reasons to avoid agony?[204] Even if Markovits could develop answers to these questions, the more basic point would remain: no matter how "controversial" some requirement seems, it is the fact that a requirement involves brute and mind-independent normative facts, *at all*, that makes it a form of realism – and such a view would therefore need to assent to the same objections that realism faces.

Conclusion

Markovits offers an important and distinctive account of Kantian morality, one that is meant to be internalist, in contrast to externalist accounts like Parfit's.

[204] See Parfit (2011a: 73-82).

However, I have argued that Markovits' account must reckon with two important concerns: first, how its requirement of 'understandability to others' can count as a procedural requirement of rationality, on Markovits' own terms; and second, whether it is simply a version of realism, since it commits to the existence of irreducibly normative facts.

10 Street's Buddhist Approach (2016)

10.1 Introduction

Among attempts to reach the 'holy grail' of meta-ethics, Sharon Street's recent line of argument is one of the most novel and fascinating: it reaches into the insights and traditions of Buddhism, suggesting that it has the resources to provide a constitutivist argument for ethical objectivity. In Street's own words, in searching for meta-ethical answers, "we should look to the Buddha and try to start thinking outside the box" (2016a: 185). Street's aim is to show that "the Buddha was right" – insofar as "we can figure out how to give his insight clear, secular analytic expression" (2016a: 187).[205]

This Chapter offers a constructive examination of Street's account. Because Street's account is at a self-professedly tentative stage of development – she outlines a number of premises but only defends some of them at length – my aim will not be to offer a complete assessment of her view, but rather to provide some constructive comment on its proposed premises and structure of argument.[206] In other words, I wish to offer some observations on what challenges *would* need to be overcome in order for the view to succeed. In this way, this Chapter can be read as suggesting *potential* concerns to Street's promised account, but in a constructive fashion towards its eventual development. I proceed in two parts: I start by explicating Street's proposed argument (§8.2), and I then raise three comments for the outlined view (§8.3).

[205] Street tempers these claims with the qualification that her project is not aimed at *interpretive accuracy* of any particular tradition: her aim "is not to be loyal to, nor to propose an accurate interpretation of, any particular tradition, Buddhist or otherwise" (2016a: 187n31). Rather, it tries to take inspiration from certain Buddhist traditions and insights, turning those into an original meta-ethical argument that can stand on its own terms.

[206] Street emphasizes the project's tentative state of development at numerous points: for instance, she writes that her paper is aimed only at making some "small inroads" in "sketching" out a view, "not to give anything remotely approaching a full answer" (2016a: 163); and she closes the paper by writing that "the view as a whole requires much more development before it is remotely clear or convincing" (2016a: 187).

10.2 Street's Proposal

Street, one of the most prominent critics of realism,[207] proposes a vindication of ethical objectivity that is in line with our mind-given presuppositions: she sets aside all claims concerning brute and mind-independent normative facts. In Street's own words, "*realism* about normative reasons is false" (2016a: 165).[208] Instead, she offers the strategy of vindicating ethical objectivity through what she calls meta-ethical *constructivism*:

> According to *metaethical constructivism*, the objectivity of ethics would be vindicated, without metaphysical and epistemological mystery, if we could somehow show that there is a universal problem, faced by every individual agent, to which morality is the universal and best (or only) solution. (Street 2016a: 166)

On such a strategy, no assertions about brute or mind-independent normative facts are ever required; rather, ethical requirements are seen as *constitutive* of how an agent can most successfully deal with some practical problem that she already faces from within her practical standpoint. Since the agent has a need to solve her own practical problem – she hence *gives herself* normative reasons to find a solution – we can say that, insofar as the *best* solution just is an *ethical* one, then she will have constitutively given herself normative reasons that have an ethical nature to them.[209]

---

[207] See especially Street (2006).

[208] Although Street shares our mind-given or antirealist presuppositions (see Chapter 1), we should note that we have not claimed that realism is *false*, but instead have found sufficient reasons to explore whether there are compelling alternatives to it (§1.2).

[209] It is worth noting the shift in Street's philosophical project over time. In earlier works, Street expressed a pessimism about the prospects of constructivist objectivity, concluding that Caligula indeed had reasons to torture others (see 2010: 371, as well as 2012). Here is how Street describes the shift: "In spite of having defended this rather subjectivist view, I don't think we should give up on the search for a stronger form of moral objectivity than Humean constructivism delivers… it would be premature to give up on the search for this Holy Grail" (2016a: 165).

The particular argument that Street offers is inspired by Buddhism, in that both the *practical problem* and *best solution* she identifies can be traced back to Buddhist traditions. The *practical problem* that applies to all valuing agents is what Street calls the *problem of attachment and loss*. According to Street, this is a *formal* problem that attaches to the nature of any act of valuing by an agent: Street claims that all agential valuing involves acknowledging either a *gap* or the *potential for a gap* between how the world is and how one would like it to be; and as a result, any valuing agent faces the problem of how to either *close this gap* or to *ensure that the potential for a gap does not open*. Here is how Street puts the problem:

> The problem that is necessarily built into the standpoint of any valuer is the problem of a gap, or the potential for a gap, between the world as it as a matter of fact is and the world as the valuer thinks it would be good or desirable for the world to be. Let us call this the *problem of attachment and loss*, with the intuitive idea being that to be a valuer is, among other things, to be a being who is attached to the world's being one way rather than another. To be a valuer, in other words, is to be a being for whom there is always *at least a potential divergence* between how things are and how one values their being, such that one is subject to loss, or at least the ever-present threat of loss. (Street 2016a: 180, emphasis added)

Street's idea is intuitive: if you care about something, then you will either face the practical problem of how to get that thing, or if you already have that thing, then you will face the practical problem of how to prevent its eventual loss. In illustrating this point, Street often appeals to the suggestion that we human beings will eventually have to confront the inevitable loss of all our loved ones and other objects of attachment:

> Each of us faces the eventual loss of everyone and everything we love. Is there a way to live in full awareness of that fact without falling into anxiety or

> depression, or resorting to one form or another of forgetfulness, denial or numbing out? (Street 2016a: 161)

This particular point is contingently true in a widespread manner for humans: most, if not all, humans do in fact care about people and other things that are vulnerable to eventual loss as time moves on. But Street is careful to note that not *every* valuer must in principle care about certain specific substantive ends, such as the existence of other people, and so the problem of attachment and loss must take on a more general framing.[210] Street is therefore explicit in formulating the problem of attachment and loss not as a problem that attaches to specific ends, but rather as a *formal* problem that attaches to any kind of valuing from an agent:

> The problem of attachment and loss is the *form* of a problem, not a problem itself. For a real problem to exist, you need a real, living being with an actual evaluative point of view on the world, a point of view constituted by actual, substantive values. Noting the formal nature of the problem of attachment and loss, we may now see the following. The 'problem of attachment and loss' – which doesn't exist except as substantively manifested in actual, contentful evaluative points of view – just inherits all of the urgency of whatever problems a given agent in fact faces… If your most urgent problem is that you're about to starve to death, then that is the force of the problem of attachment and loss for you. If you are anguished about the inevitable loss of everyone and everything you love, then that is the force of the problem of attachment and loss for you… it is just a formal description of *whatever* is the most urgent problem the agent faces. (Street 2016a: 183-184)

This problem therefore applies to anyone who values at least *something* that is not intrinsically permanent and continually secure, and the problem therefore

[210] For instance, Street imagines valuers who do not value loved ones who might pass on, but instead, unchanging facts: "a valuer who values the fact that a certain book was written, or that a certain piece of artwork was produced, or that a certain love affair was had – *and nothing else*" (2016a: 181n26, emphasis added).

applies to nearly all imaginable human beings, as well as all valuing *agents* – all valuers who have some reason to *do* something in virtue of their valuing.[211] It is also important to notice that Street emphasizes that the *urgency* of the problem just takes on the same urgency of our strongest normative concern, which is what allows her to make the claim that coming to a solution for the gap is, at a definitional or analytic level, necessarily our *most urgent* practical problem.[212] And the problem of attachment and loss can be seen in its Buddhist origins, insofar as the historical Buddha spoke to a rather deep problem about "desire and impermanence" – a problem that, in Street's terms, is about seeing a gap between how the world is and how one would like it to be.[213]

What is Street's proposed solution to the problem of attachment and loss? Street suggests that the *best* (or *only*) solution to this practical problem involves reaching a state of mind that she calls the *standpoint of pure awareness*. The standpoint of pure awareness involves a number of characteristics, foremost of which is to see the world from a maximally thin and *universal* point of view:

> Arguably, there is a maximally thin, universal point of view on the world that one may reach via the practice of mindfulness meditation. Call this point of view the *standpoint of pure awareness*… when one occupies this standpoint, one is identifying with a universal point of view on the world that transcends

[211] Street's argument therefore does not cover what she calls *invulnerable valuers*: those who are omnipotent, and those who value unchanging facts, such as the fact that two plus two equals four (Street 2016a: 182). Street is right to note that such people are uncommon, but she also offers a deeper reason for not being concerned about them: people who value facts like two plus two equals four have no reasons *for action*. They are therefore not the pressing objects of our theoretical concern, since they have no reasons to do anything at all – hence Street targets not valuers *qua valuers* in her argument, but rather *valuing agents*, i.e. those who value and have reasons for action.

[212] Indeed, given that the practical problem of attachment and loss is just the *form* of a problem that attaches to any act of substantive valuing, it seems that any ordinary person will technically have multiple such problems, from wanting to eat chocolate to not wanting to lose all of one's loved ones. Street is therefore right to focus on the *strongest* of these normative concerns as the most urgent practical problem an agent faces.

[213] For instance, Street writes that "the historical person of the Buddha, after all, may be understood has having been speaking directly to the problem of attachment and loss (pitching it as a problem about desire and impermanence), and he suggested that there is a universal solution" (2016a: 185).

> any particular, finite point of view; in doing so, one is identifying with a point of view that is not itself vulnerable to loss. (Street 2016a: 186)

Street suggests that we can reach this "state of mind" through the practice of mindfulness meditation: "a form of mental training in which one cultivates the ability to pay attention, directly and non-judgmentally, to everything that is happening in one's field of awareness in the present moment" (2016a: 185-186).[214] And Street claims that this is a path that "the historical person of the Buddha… recommended" (2016a: 185-186).

There are two important things to note about the standpoint of pure awareness. The first is that it is supposed to constitute a *solution* to the problem of attachment and loss – indeed, it is supposed to constitute the *best* or *only* solution to the problem of attachment and loss for any given valuing agent. To illustrate this, Street tells the story of Kisa Gotami from the Buddhist tradition, who loses her only treasured son, and is then in significant grief. Having lost her son, Kisa Gotami faces the problem of attachment and loss: "Kisa Gotami's son was dead and she would never get him back. As individuals, we are vulnerable to unmitigated loss" (2016a: 185). Then, here is how Street describes how reaching the standpoint of pure awareness can constitute a *solution* to her problem:

> Arguably, the standpoint of pure awareness offers each individual valuer the solution to his or her own particular version of the problem of attachment and loss... Kisa Gotami does not get her son back, but she *comes to identify with a maximally universal point of view.* This does not change her loves so much as her relationship *with* them. She still adores and misses her son, but she sees that 'she' is not ultimately the same as these loves – that these loves are rather part of the world like everything else, and that in the end 'she' and her son and every other being are the same. (Street 2016a: 186, emphasis added)

[214] Street does not claim that meditation is the *only* way to reach the standpoint, but it is "currently the most reliable means we have of reaching that state of mind" (2016a: 186).

According to Street's conception, it is not that Kisa Gotami no longer misses her son – she does still miss him. But rather, she comes to *identify* with a maximally universal point of view, which changes her *relationship* with her love for her son. Since she is no longer identifying with her own love *as hers*, we might describe her as coming to see the love in terms of love *qua* love, but not something that she *owns* or *possesses*.[215] Of course, more can be said here about the nature of this solution, and crucially, how it can constitute the *best* or *only* solution to the problem – topics that I shall return to in the next section. The point at present is that the standpoint is envisaged, in some way, as transforming the nature of problem of attachment and loss.

The second important point about the standpoint of pure awareness is that it is also supposed to be an *ethical* standpoint – one in which the agent has ethical reasons for action – and this is what brings the argument into its constitutivist fruition, as a vindication of ethical objectivity. For Street claims that the standpoint of awareness is one in which the agent has compassion for all beings:

> Arguably, the standpoint of pure awareness is also an ethical standpoint. Why? Because the point of view in question is that of pure subject – the maximally thin point of view of the 'one who is aware'. The point of view sees itself everywhere and is necessarily compassionate with the lived experience of every being. It recognizes that there is no real difference between 'self' and 'other' – that instead 'I' am everywhere, in 'you' as much as 'me'. (Street 2016a: 186-187)

Since Street's general meta-ethical view is that our normative reasons are grounded in our normative attitudes, given that someone who has reached the standpoint of pure awareness is in a state of mind of having compassion for all beings, that person will presumably also have normative reasons to care and act

[215] This is how I read Street's suggestion that "these loves are rather *part of the world* like everything else" (2016a: 186, emphasis added).

in light of that compassion.[216] So, insofar as the standpoint of pure awareness involves this kind of universal compassion, it promises to be an ethical standpoint that provides us ethical reasons for action. Of course, more can be said about this claim as well, which we will turn to in the next section.

One can therefore see the intended arc of Street's argument: from the fact that we all face normative problems that can be described under the problem of attachment and loss, we have decisive reasons to find a best solution to those problems; and, insofar as the best solution is to take up an ethical perspective on the world, it can be shown that we have decisive reasons to take up an ethical perspective on the world. Here is the argument, broken into four premises:

1. Every valuing agent, in virtue of her valuing, faces the practical problem that there is a mismatch or potential mismatch between how the world is and how she would like the world to be – the *problem of attachment and loss*. She therefore has normative reasons to find a solution to this problem.

2. One of these mismatches or potential mismatches will constitute the most urgent practical problem that the valuing agent faces, and her reasons to solve it will therefore be her most decisive reasons.

3. The best (or only) solution to this practical problem is to reach a maximally thin and universal standpoint on the world – the *standpoint of pure awareness*.

4. The standpoint of pure awareness is also an ethical perspective, in which one has ethical reasons that concern all beings.

5. Conclusion: Every valuing agent, in virtue of her valuing, has decisive reasons to reach an ethical perspective, in which she has ethical reasons for action. This would constitute a vindication of the objectivity of ethics.

[216] See Street (2008, 2010, and 2017) for her general view on how normative reasons are grounded in normative attitudes.

In her paper, Street focuses most of her efforts on developing and defending the first two premises, while she leaves the third and fourth premises at an introductory level, for further defense at a later point.[217]

This is an argument that demands our attention in virtue of its ingenuity and inventiveness, as well as its refreshing ambition of drawing upon the resources of Buddhism. In the following section, I will raise three comments on this view in development: one about the nature of the standpoint of pure awareness, another on the *structure* of the argument, and a third about what it would take for the standpoint of pure awareness to be the *best* solution to the problem.

10.3 Three Comments

This section raises three comments on Street's argument. Let me describe these in brief, before developing each in turn. The first comment concerns the nature of the standpoint of pure awareness itself. In order for Street's argument to work, the standpoint of pure awareness must simultaneously serve two roles: it must at once solve the practical problem of attachment and loss, while also remaining a source of ethical reasons *for action*; and I shall suggest that, at least at face value, there is a tension here. Second, I shall suggest that there is a structural concern with the argument: even if all of its premises are successful, the argument seems to suggest only that all valuing agents have reasons to *come to acquire* ethical reasons for action, but not that they actually have ethical reasons for action. Third, I shall also raise some questions about what it would take for the standpoint of pure awareness to constitute a *best* solution to the problem of attachment and loss, including raising some notes on alternative options related to theism.

[217] Street claims to provide only a "hasty sketch" of the third premise, with the aim, in this paper, of showing how the third and fourth premises *might* be true (2016a: 175). As I understand it, Street is now developing some of these ideas into a book manuscript.

First, let me raise a question about the nature of the standpoint of pure awareness itself. Street only describes the nature of the standpoint in brief in her paper, limited mainly to the passages I have already quoted, and so my commentary on it will be to some extent based on speculative construction. The presumptive tension I want to point out goes as follows. On the one hand, the standpoint of pure awareness must involve enough detachment from our practical concerns such that it can actually constitute a *solution* to the problem of attachment and loss – we must in some way come to no longer feel attached, moved, or bothered by the fact that we will lose our loved ones and other objects of attachment as time moves on, or the like. But on the other hand, the standpoint of pure awareness must not involve too much detachment such that we are no longer *moved to action*, since it is still supposed to be a wellspring of ethical reasons *for action* – we are supposed to take on compassion for all beings, but then to see their suffering as things that count in favor of our *doing something*. So the question is how both of these things might be achieved, in a single conception of a "state of mind" that we can be in.

We can find suggestions from Street as to how both dimensions of the tension can rendered simultaneously compatible. For instance, let me quote again a passage that touches on these issues:

> Arguably, the standpoint of pure awareness offers each individual valuer the solution to his or her own particular version of the problem of attachment and loss. Not in one sense, but in another. How so? Because when one occupies this standpoint, one is identifying with a universal point of view on the world that transcends any particular, finite point of view; in doing so, one is identifying with a point of view that is not itself vulnerable to loss. Kisa Gotami does not get her son back, but she comes to identify with a maximally universal point of view. This does not change her loves so much as her relationship *with* them. She still adores and misses her son, but she sees that 'she' is not ultimately the same as these loves – that these loves are rather part of the world like everything else, and that in the end 'she' and her son and every other being are the same. (Street 2016a: 186)

The suggestion that emerges from this passage is as follows: when we take on the standpoint of pure awareness, we still *care* about things in the world, but we do not care about things in a way that makes an essential reference to our *own* attachment to them. For instance, Kisa Gotami still loves her son, but this love is no longer attached to her own sense of self – rather, the love is "part of the world like everything else." This seems, at least in part, to be what makes the standpoint a *universal* point of view – in line with the Buddhist tradition, it is one in which one's sense of self dissipates, and the 'agent' just has universal concern for all sentient sources of suffering in the world.

This would seem to offer a resolution to the tension I described, in that (i) the 'agent' still *cares* about things, in a way that still provides her space to have *reasons for action*; (ii) but at the same time, because the agent *herself* is not attached to particular things in a *possessive* or essentially *self-referencing* manner, the standpoint also constitutes a solution to her original practical problem. But this proposed resolution then runs us into another question, which concerns (ii) – and this has to do with how the practical problem of attachment and loss was originally formulated. Recall that, when the problem of attachment and loss was introduced by Street, it was a *formal* problem that attached to any act of valuing:

> The problem that is necessarily built into the standpoint of any valuer is the problem of a gap, or the potential for a gap, between the world as it as a matter of fact is and the world as the valuer thinks it would be good or desirable for the world to be. Let us call this the *problem of attachment and loss*, with the intuitive idea being that to be a valuer is, among other things, to be a being who is attached to the world's being one way rather than another… it is just a formal description of *whatever* is the most urgent problem the agent faces. (Street 2016a: 180, 184)

Now, on Street's conception of the enlightened Kisa Gotami – the version of Kisa Gotami who has reached the standpoint of awareness – she still loves her

son, and indeed sees *everybody's* suffering in the world as constituting normative reasons for action. If so, then there is still very much a *mismatch* between how the world is and how the Kisa Gotami would like the world to be, in virtue of her universal compassion: in fact, this mismatch seems to *enlarge* in an extremely significant way, since now Kisa Gotami is aware of and compassionate for *all* the suffering in the world, not just the concerns that were previously once limited to her own set of attachments. So, the question we are now entitled to raise is the following: how has the problem of attachment and been solved, rather than enlarged in its scope and depth?

It seems that the proposed picture cannot maintain that there is *no longer a mismatch*, since the nature of claiming that there are still ethical *reasons for action* from the standpoint involves implying that there remains a mismatch between how the world is and how one would like it to be, with normative implications.[218] What Street might want to claim is that the *nature* of the problem has transformed, since while there is still normative concern for things from within the standpoint of pure awareness, these are no longer *Kisa Gotami's concerns*, in a possessive sense that makes essential reference to the self. If that is the line that is taken, then the problem of attachment and loss would need to be amended in its original formulation as follows: it is the problem of a mismatch, or potential mismatch, arising between how the world is and how one would like the world to be, *as a problem related to an individual's sense of self being attached to this mismatch*. But notice that, at least in Street's explicit formulation of the problem, this is not how it was originally stated:

> The problem that is necessarily built into the standpoint of any valuer is the problem of a gap, or the potential for a gap, between the world as it as a matter of fact is and the world as the valuer thinks it would be good or desirable for the world to be. (Street 2016a: 180)

[218] We might even put the point as follows: that it is part of the very *concept* of having a *normative reason for action* that there is an implied mismatch or potential mismatch between how the world is and how one would like it to be.

This way of formulating the problem puts it in general terms, without making reference to whether an individual's sense of self is attached to the mismatch or potential mismatch. So the problem of attachment and loss, at least as it was originally formulated, would still retain full force from within the standpoint of pure awareness: for even from that practical standpoint, the kind of mismatch or potential mismatch that Street describes would still exist.

The first comment can therefore be summarized as follows: in order for the eventual view to succeed, it must describe the precise sense in which the practical problem of attachment and loss is *solved* upon reaching the standpoint of pure awareness, while still acknowledging a mismatch or potential mismatch that would move an agent to adopting ethical reasons for *action*.[219]

Second, I want to raise a structural concern about the proposed picture: even if each of its premises are successful, what it seems to show is that we have reasons to *come to acquire* ethical reasons, but not that most of us *actually have* ethical reasons. This is because the theory seems to show that we have reasons to attain the standpoint of pure awareness, but prior to attaining it, we do not yet have ethical reasons for action. This is due to the following observation: the point at which we acquire ethical reasons, according to the argument, is *after* we have reached the standpoint of pure awareness, since that is the state of mind from which we have compassion for all beings. Prior to reaching that standpoint, the argument seems to show that we have what might be called *proleptic* reasons to work towards achieving that state of mind, since we have reasons to seek the best solution to our practical problem.[220] But this would not show that we *now* have reasons to help those around us. So, the structural concern could be put as

[219] One suggestion here might be that the agent from the standpoint of pure awareness just *does act*, without thinking about it or acting *on a reason*. The agent just sees suffering, and almost like some kind of muscle reflex, acts upon the suffering in a compassionate manner. This kind of proposal would remove the sense in which the agent must be *aware* of a mismatch or potential mismatch, in some kind of conscious manner, though it would raise other questions. Another suggestion might be that, while the agent is aware that certain actions are called for by certain situations, she does not end up feeling devastated or concerned if her actions are unsuccessful in achieving their aims – and in this sense, she is no longer concerned with the problem of attachment and loss.

[220] I borrow the notion of a *proleptic* reason from Callard (2018) – it is one that is concerned with aspiring towards being in a certain kind of state.

follows: even if each premise were vindicated as successful, the argument would show that Hitler had normative reasons to later acquire ethical reasons, through reaching the standpoint of pure awareness at some point – but it would not show that the *actual Hitler*, who clearly had not reached the standpoint of pure awareness, had normative reasons to stop committing the Holocaust.

Here is why this seems to be the implication of the structure of Street's proposed argument. On the most natural reading of Street, her meta-ethical constructivism entails that an agent's normative reasons for action are grounded in her *actual* practical standpoint, in combination with the logical and instrumental facts – rather than some *hypothetical* or *future* practical standpoint.[221] I take it that Street's argument is intended to show that, no matter what one's *actual* practical standpoint is, one has decisive reasons to work towards attaining the standpoint of pure awareness – since the problem of attachment and loss exists in all actual practical standpoints of valuing agents. Crucially, however, the actual practical standpoints of these valuing agents is not *yet* one that involves compassion for all beings: and since our reasons are grounded in our actual practical standpoints, such agents do not yet have reasons to care about and help all those who are suffering. What Hitler has normative reason to do, according to the argument, is to spend a lot of time meditating[222] – but until that meditation bears fruit, he has not yet acquired the compassionate state of mind that would then entail a reason to stop the Holocaust from occurring.

One way Street could resist this point is by suggesting that our hypothetical or future practical standpoints do ground normative reasons for us in the present. On such a suggestion, insofar as we have reasons to attain some possible practical standpoint P that involves having ethical reasons E, then we

[221] For instance, this seems to be the upshot of Street's discussion about ideally coherent eccentrics (2009a). When Street discusses someone who does not care about his own future pain, she writes: "If we imagine someone who in full coherence and making no non-normative mistakes whatsoever shows absolutely no regard for the interests of his future self at a given time or set of times, then we may discover that the issues in play are actually more ones of morality than prudence" (2009a: 292). The implication is that there is not some prudential requirement of rationality to care about the future, insofar as one just does not currently care.
[222] Or some other means of reaching the standpoint of pure awareness.

just do *now* have ethical reasons E. But we might find that this suggestion faces difficulties, for the following reasons. First, it does not seem to be Street's position, as far as her previous published works are concerned.[223] Second, the view would seem to have counterintuitive implications on its own terms, when taken as a general principle: for instance, if I have a reason to take a shower after a long day of work in the fields, that does not seem to entail a reason to take a shower before going into the fields. Moreover, the view would need to be motivated on constitutivist or antirealist grounds, in order for it to align with the presuppositions of Street's meta-ethical position. Finally, note that even on recent subjectivist suggestions that our futures are related to our present reasons, the view is not that the reasons we will have just are the reasons we now have: rather, the view is often that our future *well-being* constitutes present reasons, which is different from the proposal now under consideration.[224] It therefore seems that it would not be straightforward to claim that the reasons we will have, from some future or hypothetical standpoint, just are the reasons we now have, such that the ethical reasons we have from the standpoint of pure awareness just are the ethical reasons we have now.[225]

A further point that could be made in response to this structural concern is that we need not vindicate that people have ethical reasons in their actual, unenlightened forms – it is theoretically *satisfactory* to show that these people have normative reasons to later acquire ethical reasons for action. The first thing that should be said, to such a response, is that demonstrating that Hitler had reasons to come to acquire ethical reasons, through meditation or otherwise, would be a significant and notable meta-ethical achievement – it would be non-trivial to have vindicated that conclusion. But the second thing to be said is that, still, this would not what would capture the full nature of the objectivist intuition, which is the thought that the actual Hitler, who set into motion the Holocaust,

[223] For instance, such a suggestion is not found in the 2016a paper, and it does not seem to be the most natural reading of her previous work (e.g. 2006, 2008, 2009a, 2010, 2012).
[224] For instance, see Sobel (2011) and Lin (2020).
[225] Street could also attempt to argue that the best solution to the problem of attachment and loss just is to directly take ethical actions, though this does not seem to be the suggestion in the paper at hand.

had decisive and serious ethical reasons to never have done that. What we want to say is that the actual Hitler made an egregious normative mistake, not only in failing to meditate and later come to acquire ethical reasons, but in failing to follow the ethical reasons that he already had – reasons not to have harmed and killed countless numbers of people. Indeed, more generally, what we want to vindicate most deeply is that valuing agents, *now*, have reasons to be concerned about the suffering and well-being of others, and to *do* things in virtue of that.[226]

This point is especially salient, because we might imagine attaining the standpoint of pure awareness is *hard*. Moreover, we might imagine that, given that the standpoint of pure awareness is a "state of mind," it might furthermore be hard to *maintain* it, after having attained it. For it seems plausible to assume that not *many* of us have actually ever achieved a maximally thin and universal point of view that "sees itself everywhere and is necessarily compassionate with the lived experience of every being… [with] no real difference between 'self' and 'other'" (Street 2016a: 186-187). Indeed, one of the points that seems to emerge from Buddhist tradition and experience is that the attainment of such a point of view may require years, if not a lifetime, of hard work: some monks dedicate their entire livelihoods to this task, and not always to consistent success.[227] But if a theory cannot maintain that we have ethical reasons *until* we reach such a standpoint – and such a standpoint is hard to actually attain, for most human beings, in practice – then the theory will face questions about how many actual human beings can actually come to have the constitutive ethical reasons in question.

So, we can summarize the second, structural concern to Street's argument as follows: that even if all of its premises are sound, it will not have

[226] Another suggestion is that Street could claim that, if Hitler has decisive reasons to meditate, then at least he would not be able to actively spend his time harming others. In this way, having decisive reasons to pursue the standpoint of pure awareness could serve as a 'temporary block' on unethical action. Still, we might wonder whether valuing agents have *positive* ethical reasons to help others, prior to reaching the standpoint of pure awareness.

[227] We might also wonder how we are supposed to maintain being in the standpoint of pure awareness, when we are then meant to take actual *actions* in the real world – something which might disrupt the balance of a certain tranquil state of mind, especially if the ethical action in question is demanding, like making sustained efforts to save someone's life.

shown that most actual human beings have ethical reasons for action – or would ever come to constitutively acquire these ethical reasons for action in a stable manner in their lifetimes, given the difficulty of attaining and maintaining the standpoint of pure awareness. What the currently formulated view seems to demonstrate, granting that its premises are sound, is that we all have *proleptic* reasons to aspire towards a state in which we then have ethical reasons for action.

Third, and most briefly, I want to make some speculative comments about what it would take for the standpoint of pure awareness to constitute the *best* solution to the problem of attachment and loss. Street often notes that, strictly speaking, what the account must show is that the standpoint of pure awareness must constitute the "*best* (or *only*) solution" to the problem (2016a: 166, 174, 175, emphasis edited). However, it seems that even Street concedes that the standpoint of pure awareness is unlikely to be the *only* solution to the problem: for we could always change *what* we value to things that are invulnerable to loss, or seek to limit our aims to things that are achievable through action, or resort to forgetfulness of the problem (Street 2016a: 162-163), or adopt some form of theism (Street 2016a: 163). It therefore seems that the most plausible approach is not for Street to claim that the standpoint constitutes the *only* solution to the problem, but rather that it constitutes the *best* solution to the problem – i.e. a solution that is *better* than all alternatives.

Obviously, a full discussion of these comparative approaches would be a large task – a point that Street recognizes.[228] I will therefore not aim in these final remarks to comment on that overall question. Rather, what I want to do is to raise some speculative suggestions on how *theism*, in particular, might fare as an alternative solution to the problem, especially it is an option that Street only gives limited consideration to in the paper.[229]

---

[228] Street writes that showing that the standpoint of pure awareness is the *best* solution is "the most puzzling and formidable" of her entire argument, and that "it is not the point of this paper to explain or defend this step" (2016a: 184).

[229] Street only briefly mentions the possibility of theism as a solution to the problem, writing that "theistic belief offers solace to some, but if one finds it impossible to believe in the existence of an omniscient, omnipotent, and perfectly good god, then that is not an option" (2016a: 163).

There is a distinguished philosophical tradition of recognizing *practical reasons* for faith. For instance, this seems to have been Kant's position, when he argued that there were practical reasons to have faith in propositions concerning immortality, God, and freedom. Among these were reasons to find solutions to what we could describe as practical problems: for instance, the practical problem of being able to sustain ourselves, motivationally, in doing the right thing, given that we have an expectation that happiness and morality should ultimately find some kind of alignment, which we do not necessarily see in the real world.[230] Kant's view, of course, was premised on the idea that there are certain theistic propositions that cannot be known through *knowledge*, and that is what makes them permissible objects of practical faith – and the thought here might be that there are certain propositions that cannot be either proven or disproven, but which could help with assuaging the problem of attachment and loss, if they were taken on faith. More recently, Lara Buchak has also developed a view in which, in light of epistemic uncertainty, there can be practical reasons for faith, indeed ones which may even be *rationally required*:

> Faith… is rationally required in circumstances in which the costs of delaying the decision are high enough to outweigh the benefit of additional evidence… in particular, faith is more apt to be rational if potential evidence against *X* will be inconclusive. (Buchak 2013: 538)

The point here is not to adopt the specific proposals of either Kant or Buchak's view, but to suggest that, in general, there might be practical reasons for taking on certain propositions *on faith* – reasons that stem from what we need in practical life, in response to practical problems.

---

[230] See Kant in the first *Critique* for his famous claim that he had to deny knowledge in order to create room for faith (Kant C1 BXXX and B856-857). He then discusses the immortality of the soul as an object of practical faith in the second *Critique*, as something it is "necessary to assume" (C2 5:122). He makes comparable remarks about the idea that we live in the best of all possible worlds, writing that "even our *practical* reason takes a great interest in this doctrine and knows it to be a necessary presupposition for its own sake and without founding it on theology" (Kant 1978: 139).

How might this be relevant in considering alternatives to Street's proposed solution? In order for Street to be able to vindicate her claim that the standpoint of pure awareness is the *best* solution to the problem of attachment and loss, she would need to rule out alternative options that involve practical faith as less attractive, either in terms of their effectiveness, ease, depth of resolution, or some other relevant criterion that is implied in some solution being the 'best.' But there are at least a number of possible objects of practical faith that could conceivably constitute a solution to the problem of attachment and loss. Let me just suggest some, in a very speculative fashion. One is the proposition that, *eventually*, the mismatch or potential for mismatch between how things are and how one would like them to be will come to a close – with faith in this proposition assuaging a person's sense of inevitable loss.[231] Along these lines, another common theistic proposition is that there is a more perfect world to come. A somewhat less common line of thought might be that we live in the best of all possible worlds, which might assuage the sense that the mismatch or potential mismatch is one that we have much to lament about, apart from what we can control in our actions.[232] More broadly, there might be an even more general attitude of *trust* or *faith* that there is *some* kind of ultimate resolution or explanation for the normative shortcomings that we see in the world, even if we do not know the precise nature of that ultimate resolution or explanation. Not all of these propositions must be tied to the idea that there is an omniscient, omnipotent, and omnibenevolent God – they need not even be tied to monotheism – although some of them certainly come along quite naturally with such propositions. The point here is not to develop or defend any particular proposal, but to suggest that these alternatives that are on the table must be considered, before we can vindicate the claim that the standpoint of pure awareness is the *best* solution. And this point is particularly salient for Street's

[231] There is space to explore whether this proposition could even be understood in a non-theistic manner, as an act of practical *hope* in what humanity could attain in the future.

[232] This proposition was defended by Plato, Leibniz, and Kant (see Plato's *Timaeus* 30a-b, Leibniz's *Theodicy* §213-215, and Kant 1978: 139), though it seems to have fallen into neglect in recent philosophy of religion. I attempt to recover attention to the doctrine in my manuscript, "Reconsidering the Best of All Possible Worlds."

project, for she herself has previously defended the idea that even the normativity of epistemic reasons is mind-dependent: "whether someone has reason to be engaged in the belief business in turn depends on whether someone has *practical* reasons… and as I've argued elsewhere, whether we have any *practical* reasons ultimately depends on whether we take ourselves to have them" (Street 2009b: 246). The suggestion here would be that belief in certain theistic propositions could similarly be guided by such practical reasons, towards a resolution to the problem of attachment and loss.

While faith in certain theistic propositions based on practical reasons might invite their own philosophical objections, with requisite discussion needed, there are also some conceivable advantages to such practical faith, in comparison to aspiring towards the standpoint of pure awareness. First, there is space to argue that such practical faith is more *immediate* and *achievable* than reaching the standpoint of pure awareness – while the adoption of a mindset of faith might not happen in a matter of hours or days, it does seem to be more directly influenced by our phenomenological efforts, in comparison to seeking the results of long meditation practice towards a maximally thin and universal state of mind. Second, there is also room to argue that such practical faith could prove more *stable* than the standpoint of pure awareness, i.e. comparatively easier to maintain, insofar as universal compassion is a difficult state of mind to always keep oneself in.[233] And third, as an empirical point, it seems that many people *have* achieved such forms of faith in real life, given the prevalent existence of religious belief, whereas there seems to be more room to question whether the envisioned point of enlightenment from the Buddhist tradition is indeed as widespread in actual attainment, even though it certainly is a widespread fact that many people are *striving* towards it.

This is not to say that there are not also comparative benefits to the solution of the standpoint of pure awareness – for instance, as Street envisions it, such a standpoint involves no need for further metaphysical or theological

[233] This is not to say that people do not undergo crises of faith – it is merely to sketch out space for a comparative point, *relative to* the ease of maintaining the standpoint of pure awareness.

beliefs which are vulnerable to their own set of objections. The point at hand is to simply suggest that, in order for Street's view to be vindicated, it would need to show that the standpoint is decisively *better* than other forms of solutions, including the suggestions here about theism, alongside other conceivable alternatives.

Conclusion

Street's proposal of a Buddhist-inspired constructivism is one of the most exciting and welcome contributions to the recent literature, in searching for the 'holy grail' of meta-ethics – a mind-given vindication of the objectivity of ethics. In this Chapter, I have explicated Street's proposed picture, and I have raised three constructive points in response to it: (i) there is a possible tension in how the standpoint of pure awareness is explicated; (ii) there is a structural problem in how the theory seems to vindicate reasons to *become* ethical, rather than to *be* ethical; and (iii) we can wonder whether theism might be a competing or comparatively advantageous practical solution to the practical problem. These points should be taken as constructive comments towards a view in development, rather than as objections towards a developed view. And, given that there is much that could be said in response to these concerns, there is much to anticipate in this novel and refreshing meta-ethical position.

11 Conclusion

This dissertation has presented a theory of normativity, along with an examination of recent work in meta-ethics. Our positive vision has been to vindicate that there are certain objective moral and other normative requirements. There are objective moral requirements because our valuing commits us to the conclusion that valuing *qua* valuing is sufficient for love, and therefore that all valuing perspectives matter; and there are other objective normative requirements because our practical standpoints involve a commitment to the normativity of a range of objective mind-world potentialities, alongside other norms of endorsement. All of these claims were presented without asserting the existence of brute and mind-independent normative facts; and I furthermore developed a picture of how our minds render the world into a normative place, and explained why a mind-given theory of normativity need not be nihilistic, for potentialities of normativity are timeless and everlasting. Finally, I examined a number of important and influential views in the literature that neighbor our own project, raising some critical and constructive points towards them.

At the close of this dissertation, I want to remind my reader that I have not claimed that there are no further normative principles or truths: for as I emphasized from the outset, this dissertation takes no stance on whether theism or realism are true. Rather, it has motivated a case for independently exploring whether we can vindicate normative objectivity without these weightier assumptions, and its closing suggestion is that we can. It has tried to show that we can find a *floor* for normativity that does not require either theism or realism: that, even if someone does not believe in God, or does not believe in the existence of normative facts that are "out there," she still cannot escape the binding force of objective moral and other normative requirements. In other words, even if we take the relativist's own metaphysical starting point, we can show that she must be committed to morality and more.

It is sometimes said that, according to realism, we must *discover* certain normative truths that exist, whereas without realism, this is not so. But on the

position that we have advanced, there is so much that we must discover: we must discover that there is a world of normative potentialities, that other beings are real, that they can experience things as good and bad, and that we cannot help but be committed to the importance of their good and bad. We cannot escape these moral demands – they are demands that arise from the nature of our conscious being, and it is to this extent that Kant was right to describe the "moral vocation of our nature" (C2 5: 122). Theism or realism might be true – but even without appealing to those further truths, we can say that human beings cannot help but bind themselves to the conclusion that all potentialities and perspectives matter.

Total Words: 74,815.

Works Cited

Anscombe, E. (1957 [2000]). *Intention*. Cambridge: Harvard University Press.

Aristotle. (4th Century BC [2014]). *Nicomachean Ethics*. R. Crisp (trans.) Cambridge University Press.

Aquinas, T. (1485 [1962]). *Summa Theologica*. Rome: Editiones Paulinae.

Barnes, E. (2017). "Realism and social structure." *Philos Stud* 174: 2417-2433.

Berker, S. (2020). "Quasi-Dependence." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 15*. Oxford University Press.

Blackburn, S. (2001). *Ethics: A Very Short Introduction*. New York: Oxford University Press.

Block, N. (1995). "On a confusion about a function of consciousness." *Behavioral and Brain Sciences* 18: 227-287.

Brandt, R. (1979). *A Theory of the Good and the Right*. Oxford University Press.

Brink, D. (1984). "Moral realism and the sceptical arguments from disagreement and queerness." *Australasian Journal of Philosophy* 62 (2): 111-125.

Bruckner, D. (2016). "Quirky Desires and Well-Being." *Journal of Ethics & Social Philosophy* 10 (2): 1-34.

Buchak, L. (2013). "Can it be Rational to have Faith?" In L. Pojman and M. Rea (eds.), *Philosophy of Religion: An Anthology, 7th Edition*. Wadsworth: 524-541.

Bukoski, M. (2016). “A Critique of Smith’s Constitutivism.” *Ethics* 127: 116-146.

Bukoski, M. (2018). “Korsgaard’s Arguments for the Value of Humanity.” *Philosophical Review* 127 (2): 197-224.

Callard, A. (2018). *Aspiration*. Oxford University Press.

Chang, R. (2002). “The Possibility of Parity.” *Ethics* 112: 659-688.

Chang, R. (2004). “Can Desires Provide Reasons for Action?” In R. Wallace, P. Pettit, S. Scheffler & M. Smith (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford University Press: 56-90.

Chang, R. (2013). “Grounding practical normativity: going hybrid.” *Philosophical Studies* 164 (1): 163-187.

Chang, R. (2016). “Comparativism: The Grounds of Rational Choice.” In E. Lord & B. Maguire (eds.), *Weighing Reasons*. Oxford University Press: 213-240.

Chang, R. (2021). “What is it to be a rational agent?” In R. Chang & K. Sylvan (eds.), *The Routledge Handbook of Practical Reason*: 95-109. New York: Routledge.

Cohon, Rachel. (2000). “The Roots of Reasons.” *Philosophical Review* 109: 63-85.

Cooper, J. (1997). “Introduction.” In J. Cooper (ed.), *Plato: Complete Works*. Indianapolis: Hackett.

Crisp, R. (1997). *Mill on Utilitarianism*. London: Routledge.

Crisp, R. (2005). "Value, reasons and the structure of justification: how to avoid passing the buck." *Analysis* 65 (1): 80-85.

Crisp, R. (2006). *Reasons and the Good.* New York: Oxford University Press.

Crisp, R. (2021). "Well-Being." E. Zalta (ed.) *Stanford Encyclopedia of Philosophy* (Winter 2021 Edition).

Dancy, J. (2000). *Practical Reality.* New York: Oxford University Press.

Darwall, S. (2003). "Desires, Reasons, and Causes." *Philosophy and Phenomenological Research* 67 (2): 436-443.

Davidson, D. (1963). "Actions, Reasons, and Causes." *The Journal of Philosophy* 60 (23): 685-700.

Descartes, R. (1637 [2006]). *A Discourse on the Method*. I. Maclean (trans.) New York: Oxford University Press.

Dorsey, D. (2012). "Subjectivism without Desire." *Philosophical Review* 121 (3): 407-442.

Dorsey, D. (2017a). "Why Should Welfare 'Fit'?" *The Philosophical Quarterly* 67 (269): 685-708.

Dorsey, D. (2017b). "Idealization and the Heart of Subjectivism." *Noûs* 51 (1): 196-217.

Dotsey, C (2018). “Great Escape: Tchaikovsky’s Violin Concerto.” *Houston Symphony*: https://houstonsymphony.org/tchaikovsky-violin-concerto/.

Dreyfus, H. (1995). “Heidegger on Gaining a Free Relation to Technology.” In A. Feenburg & A. Hannay (eds.), *Technology and the Politics of Knowledge*, Indiana University Press: 25-33.

Dworkin, R. (1996). “Objectivity and Truth: You’d Better Believe it.” *Philosophy and Public Affairs*, 25 (2): 87-139.

Dworkin, R. (2011). *Justice for Hedgehogs*. Cambridge: Harvard University Press.

Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. New York: Oxford University Press.

Enoch, D. (2021). “Constitutivism: On rabbits, hats, and holy grails.” In R. Chang & K. Sylvan (eds.), *The Routledge Handbook of Practical Reason:* 336--348. New York: Routledge.

Fan, Z. (2021). *Ruling reasons: a defense of constructivism about reasons and objectivism about welfare*. University of Oxford, BPhil Thesis.

Finnis, J. (1980 [2011]). *Natural Law and Natural Rights*. Second Edition. New York: Oxford University Press.

FitzPatrick, W. (2005). “The Practical Turn in Ethical Theory: Korsgaard’s Constructivism, Realism, and the Nature of Normativity.” *Ethics* 115: 651-91.

Frankfurt, H. (2004). *The Reasons of Love.* Princeton: Princeton University Press.

Frick, J. (2020). “Conditional Reasons and the Procreation Asymmetry.” *Philosophical Perspectives* 34: 53-87.

Fritzson, F. (2016). “Good, Good For, and Good Relative To: Relative and Relational in Value Theory.” *Philosophy* 91: 255-267.

Garner, R. (1990). “On the genuine queerness of moral properties and facts.” *Australasian Journal of Philosophy* 68 (2): 137-146.

Gert, J. (2002). “Korsgaard’s Private-Reasons Argument.” *Philosophy and Phenomenological Research* 64: 303-324.

Gewirth, A. (1978). *Reason and Morality*. Chicago: The University of Chicago Press.

Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.

Gibson, J. (1979 [2015]). *The Ecological Approach to Visual Perception.* New York: Psychology Press.

Griffin, J. (1986). *Well-Being: Its Meaning, Measurement, and Moral Importance.* Oxford: Clarendon Press.

Grotius, H. (1625 [2003]). *On the Law of War and Peace*. In F. Kelsey (trans.), J. Schneewind (ed.), *Moral Philosophy from Montaigne to Kant*. New York: Cambridge University Press.

Hare, R. (1981). *Moral Thinking: Its Levels, Method, and Point*. New York: Oxford University Press.

Hawkins, J. (2019). “Internalism and Prudential Value.” In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 14*. Oxford: Oxford University Press.

Heathwood, C. (2005). “The problem of defective desires.” *Australasian Journal of Philosophy* 83 (4): 487-504.

Hills, A. (2008). “Kantian Value Realism.” *Ratio* 221: 182-200.

Hume, D. (1739 [2007]). *A Treatise of Human Nature.* D. Norton & M. Norton (eds.) New York: Oxford University Press.

Hume, D. (1751 [1983]). *An Enquiry Concerning the Principles of Morals.* J. B. Schneewind (ed.) Indianapolis: Hackett.

Hurka, T. (1987). “’Good’ and ‘Good For’.” *Mind* 96 (381): 71-73.

Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press.

Kant, I. (1785 [1998]). *Groundwork of the Metaphysics of Morals* (G). M. Gregor (trans.) Cambridge University Press.

Kant, I. (1787 [1998]). *Critique of Pure Reason* (C1). P. Guyer & A. Wood (trans.) Cambridge University Press.

Kant, I. (1788 [2015]). *Critique of Practical Reason* (C2). M. Gregor (trans.) Cambridge University Press.

Kant, I. (1790 [2000]). *Critique of the Power of Judgment* (C3). P. Guyer & E. Matthews (trans.) Cambridge University Press.

Kant, I. (1978). *Lectures on Philosophical Theology*. A. Wood and G. Clark. (trans.) Cornell University Press.

Kerstein, S. (2001). "Korsgaard's Kantian Arguments for the Value of Humanity." *Canadian Journal of Philosophy* 31: 23-52.

Kierkegaard, S. (1846 [1941]). *Concluding Unscientific Postscript*. D. Swenson & W. Lowrie (trans.) Princeton University Press.

Kiesewetter, B. (2022). "Are all practical reasons based on value?" In *Oxford Studies in Metaethics*, R. Shafer-Landau (ed.), Volume 17: 27-53.

Kolodny, N. (2005). "Why Be Rational?" *Mind* 114 (455): 509-563.

Korsgaard, C. (1996). *The Sources of Normativity*. New York: Cambridge University Press.

Korsgaard, C. (1997). "The Normativity of Instrumental Reason." In G. Cullity & B. Gaut (eds.), *Ethics and Practical Reason*. Oxford: Clarendon Press.

Korsgaard, C. (2009). *Self-Constitution: Agency, Identity, and Integrity*. New York: Oxford University Press.

Korsgaard, C. (2018). *Fellow Creatures: Our Obligations to the Other Animals*. New York: Oxford University Press.

Korsgaard, C. (2021). "Valuing Our Humanity." In R. Dean & O. Sensen (eds.), *Respect: Philosophical Essays*. Oxford University Press: 171-191.

Langton, R. (2007). “Objective and Unconditioned Value.” *Philosophical Review* 116 (2): 157-185.

Leibniz, G. (1985). *Theodicy: Essays on the Goodness of God, the Freedom of Man, and the Origin of Evil.* E. M. Huggard (trans.) Lasalle, IL: Open Court.

Lewis, D. (1989). “Dispositional Theories of Value.” *Proceedings of the Aristotelian Society, Supplementary Volumes,* 63: 113-137.

Lin, E. (2017). “Against Welfare Subjectivism.” *Noûs* 51 (2): 354-377.

Lin, E. (2020). “Future Desires, the Agony Argument, and Subjectivism about Reasons.” *Philosophical Review* 129 (1): 95-130.

Lindeman, K. (2019). “Functional Constitutivism Misunderstood Resources: A Limited Defense of Smith’s Constitutivism.” *Ethics* 130: 79-91

Locke, J. (1689 [1979]). *An Essay Concerning Human Understanding.* Oxford: Clarendon Press.

LoLordo, A. (2021). “Locke on the guise of the good.” *Philosophical Explorations* 24 (1): 21-33.

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

Maguire, B. (2016). “The Value-Based Theory of Reasons.” *Ergo* 3 (9): 233-62.

Maguire, B. (2018). “There Are No Reasons for Affective Attitudes.” *Mind* 27 (507): 779-805.

Manne, K. (2016). “Democratizing Humeanism.” In E. Lord & B. Maguire (eds.), *Weighing Reasons*. Oxford University Press.

Markovits, J. (2014). *Moral Reason.* New York: Oxford University Press.

Markovits, J. & Walden, K. (2021). “Kantian constructivism.” In R. Chang & K. Sylvan (eds.), *The Routledge Handbook of Practical Reason:* 25-37. New York: Routledge.

McGrath, S. (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 University Press: 186-214.

Mill, J. S. (1863 [2015]). *Utilitarianism*. In M. Philp & F. Rosen (eds.), *On Liberty, Utilitarianism, and Other Essays*. New York: Oxford University Press.

Moore, A. W. (2003). *Noble in Reason, Infinite in Faculty*. London: Routledge.

Moore, G. E. (1903 [1922]). *Principia Ethica*. Cambridge University Press.

Moore, G. E. (1912 [2005]). *Ethics*. William Shaw (ed.) Oxford: Clarendon Press.

Murdoch, I. (1970 [2014]). *The Sovereignty of Good*. London: Routledge.

Nagel, T. (1970). *The Possibility of Altruism*. Princeton: Princeton University Press.

Nagel, T. (1986). *The View From Nowhere*. New York: Oxford University Press.

Nietzsche, F. (1882 [1974]). *The Gay Science*. W. Kaufmann (trans.), New York: Vintage Books

Nietzsche, F. (1873 [1976]). "On Truth and Lies in a Nonmoral Sense." In W. Kaufmann (ed. & trans.), *The Portable Nietzsche*. Viking Press.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.

O'Neill, O. (1996). *Towards Justice and Virtue: A constructive account of practical reasoning.* Cambridge: Cambridge University Press.

O'Neill, O. (2000). *Bounds of Justice*. Cambridge: Cambridge University Press.

O'Neill, O. (2003). "Constructivism in Rawls and Kant." In S. Freeman (ed.), *The Cambridge Companion to Rawls*. Cambridge University Press.

Parfit, D. (1984). *Reasons and Persons*. New York: Oxford University Press.

Parfit, D. (1997). "Reasons and Motivation." *Proceedings of the Aristotelian Society, Supplementary Volumes* 71: 99-130.

Parfit, D. (2011a). *On What Matters, Volume 1.* New York: Oxford University Press.

Parfit, D. (2011b). *On What Matters, Volume 2.* New York: Oxford University Press.

Paul, L. A. (2014). *Transformative Experience*. Oxford University Press.

Plato. (5th Century BC [1997]). *Republic*. In J. Cooper (ed.), *Plato: Complete Works*. Indianapolis: Hackett.

Plato. (5th Century BC [1997]). *Timaeus*. In J. Cooper (ed.), *Plato: Complete Works*. Indianapolis: Hackett.

Quinn, W. (1993). *Morality and Action.* Cambridge: Cambridge University Press.

Rabinowicz, W. & Rønnow-Rasmussen, T. (2000). "A Distinction in Value: Intrinsic and for Its Own Sake." *Proceedings of the Aristotelian Society* 100: 33-51.

Railton, P. (1986). "Facts and Values." *Philosophical Topics* 14 (2): 5-31.

Railton, P. (2003). *Facts and Values: Essays Toward a Morality of Consequence.* New York: Cambridge University Press.

Rawls, J. (1971). *A Theory of Justice: A Revised Edition.* Cambridge: Harvard University Press.

Rawls, J. (1993). *Political Liberalism.* New York: Columbia University Press.

Raz, J. (1999). *Engaging Reason: On the Theory of Value and Action.* New York: Oxford University Press.

Regan, D. (2002). "The Value of Rational Nature." *Ethics* 112: 267-91.

Rockwell, D. & Giles, D. C. (2009). "Being a celebrity: A Phenomenology of fame." *Journal of Phenomenological Psychology* 40 (2): 178-210.

Rosen, G. (2017). “Scanlon’s modal metaphysics.” *Canadian Journal of Philosophy* 47 (6): 856-876.

Ross, W. D. (1930 [2002]). *The Right and the Good.* P. Stratton-Lake (ed.) New York: Oxford University Press.

Scanlon, T. M. (1998). *What We Owe to Each Other.* Cambridge: Harvard University Press.

Scanlon, T. M. (2014). *Being Realistic About Reasons.* New York: Oxford University Press.

Scanlon, T. M. (2021). “Some central questions about practical reason.” In R. Chang & K. Sylvan (eds.), *The Routledge Handbook of Practical Reason*: 25-37. New York: Routledge.

Schafer, K. (2014). “Constructivism and Three Forms of Perspective-Dependence in Metaethics.” *Philosophy and Phenomenological Research* 89 (1): 68-101.

Scheffler, S. (2011). “Valuing.” In R. Wallace, R. Kumar, & S. Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon.* New York: Oxford University Press.

Scheffler, S. (2013). *Death and the Afterlife.* ed. N. Kolodny. New York: Oxford University Press.

Schroeder, M. (2007). *Slaves of the Passions*. New York: Oxford University Press.

Shafer-Landau, R. (2003). *Moral Realism: A Defence.* New York: Oxford University Press.

Shemmer, Y. (2007). "Desires as Reasons." *Philosophy and Phenomenological Research* 75 (2): 326-348.

Shemmer, Y. (2020). "Subjectivism about Future Reasons or the Guise of Caring." *Philosophy and Phenomenological Research* 101 (3): 630-648.

Shepski, L. (2008). "The vanishing argument from queerness." *Australasian Journal of Philosophy* 86 (3): 371-387.

Sidgwick, H. (1907 [1981]). *The Methods of Ethics*, Seventh Edition. Indianapolis: Hackett.

Skorupski, J. (2010). *The Domain of Reasons.* New York: Oxford University Press.

Smith, M. (1987). "The Humean Theory of Motivation." *Mind* 96 (381): 36-61.

Smith, M. (1994). *The Moral Problem*. Blackwell Publishing.

Smith, M. (2013). "A Constitutivist Theory of Reasons: Its Promise and Parts." *Law, Ethics, and Philosophy* 1: 9-30.

Smith, M. (2015). "The Magic of Constitutivism." *American Philosophical Quarterly* 52 (2): 187-200.

Smith, M. (2017). "Parfit's Mistaken Meta-Ethics." In P. Singer (ed.), *Does Anything Really Matter? Essays on Parfit on Objectivity.* Oxford University Press.

Smith, M. (2020). “The Modal Conception of Ideal Rational Agents: Objectively Ideal Not Merely Subjectively Ideal, Advisors Not Exemplars, Agentially Concerned Not Agentially Indifferent, Social Not Solitary, Self-and-Other-Regarding Not Wholly Self-Regarding.” In D. Kaspar (ed.), *Explorations in Ethics*: 59-80. Palgrave MacMillan.

Smith, M. (2021). “The State of Nature.” In M. Timmons (ed.), *Oxford Studies in Normative Ethics, Volume 11*: 271-290. Oxford University Press.

Sobel, D. (2009). “Subjectivism and Idealization.” *Ethics* 119: 336-352.

Sobel, D. (2011). “Parfit’s Case Against Subjectivism.” In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, *Volume 6*. New York: Oxford University Press.

Sobel, D. (2016). *From Valuing to Value: A Defense of Subjectivism.* New York: Oxford University Press.

Spaid, A. (2021). “Reasons Internalism and the Problem of Depression.” *Utilitas* 33: 1-16.

Stace, W. T. (1944). “Interestingness.” *Philosophy* 19 (74): 233-241.

Stern, R. (2011). “The Value of Humanity: Reflections on Korsgaard’s Transcendental Argument.” In J. Smith & P. Sullivan (eds.), *Transcendental Philosophy and Naturalism*. Oxford University Press.

Street, S. (2006). “A Darwinian Dilemma for Realist Theories of Value.” *Philosophical Studies* 127 (1): 109-166.

Street, S. (2008). “Constructivism about Reasons.” In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 3.* New York: Oxford University Press.

Street, S. (2009a). “In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters.” *Philosophical Issues* 19: 273-298.

Street, S. (2009b). “Evolution and the Normativity of Epistemic Reasons.” *Canadian Journal of Philosophy,* Supplementary Volume 35: 213-248.

Street, S. (2010). “What is Constructivism in Ethics and Metaethics?” *Philosophy Compass* 5 (5): 363-384.

Street, S. (2011). “Mind-Independence without the Mystery: Why Quasi-Realists Can’t Have it Both Ways.” *Oxford Studies in Metaethics, Volume 6.* Oxford: Oxford University Press.

Street, S. (2012). “Coming to Terms with Contingency: Humean Constructivism about Practical Reason.” In J. Lenman & Y. Shemmer (eds.), *Constructivism in Practical Philosophy:* 40-59. Oxford: Oxford University Press.

Street, S. (2016a). “Constructivism in Ethics and the Problem of Attachment and Loss.” *Aristotelian Society Supplementary Volume* XC: 161-189.

Street, S. (2016b). “Objectivity and Truth: You’d Better Rethink It.” In R. Shafer-Landau (ed.), *Oxford Studies in Meta-Ethics, Volume 11,* Oxford University Press: 293-334.

Street, S. (2017). “Nothing ‘Really’ Matters, but That’s Not What Matters.” In P. Singer (ed.), *Does Anything Really Matters?: Essays on Parfit on Objectivity*. New York: Oxford University Press.

Taylor, Charles. (1985). “What is Human Agency?” In Taylor, *Human Agency and Language: Philosophical Papers 1*, p. 15-44. Cambridge University Press.

Tchaikovsky, P. (1881). “Letter 1916 to Nadezhda von Meck.” http://en.tchaikovsky-research.net/pages/Letter_1916

Tenenbaum, S. (2013). “Guise of the Good.” In H. LaFollette (ed.), *The International Encyclopedia of Ethics*, First Edition. Blackwell.

Theunissen, N. (2018). “Must We Be Just Plain Good? On Regress Arguments for the Value of Humanity.” *Ethics* 128: 346-372.

Thomson, J. (2008). *Normativity*. Chicago and La Salle: Open Court.

Tiberius, V. (2018). *Well-Being as Value Fulfillment.* New York: Oxford University Press.

Tully, I. (2017). “Depression and the Problem of Absent Desires.” *Journal of Ethics & Social Philosophy* 11 (2): 1-15.

Velleman, D. (1992). “The Guise of the Good.” *Noûs* 26 (1): 3-26.

Velleman, D. (2009). *How We Get Along*. New York: Cambridge University Press.

Wallace, R. J. (2012). “Constructivism About Normativity: Some Pitfalls.” In J. Lenman & Y. Shemmer (eds.), *Constructivism in Practical Philosophy*. Oxford: Oxford University Press.

Williams, B. (1981). “Internal and external reasons.” In B. Williams, *Moral Luck*: 101-113. Cambridge University Press.

Williams, B. (1985 [2011]). *Ethics and the Limits of Philosophy*. Routledge.

Williams, B. (1995). “Internal reasons and the obscurity of blame.” In B. Williams, *Making Sense of Humanity*: 35-45. Cambridge University Press.

Williams, B. (2006). “The Human Prejudice.” In A. Moore (ed.), *Philosophy as a Humanistic Discipline.* Princeton: Princeton University Press.

Wolf, S. (2010). *Meaning in Life and Why It Matters*. Princeton: Princeton University Press.

Zhang, K. (2019). *On Moral Objectivity: Euthyphro, Evolution, and the Source of Universal Dignity*. University Archives, Department of Special Collections, Princeton University Library.

Zhang, K. (2021). *Prolegomena to an Internalist Meta-Ethics*. University of Oxford, BPhil Thesis.

Zhang, K. (Manuscript). “On Kant’s Meta-Ethics.”

Zhang, K. (Manuscript). “Reconsidering the Best of All Possible Worlds.”