

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

UNIVERSITY OF
Southampton

Faculty of Humanities

Department of Philosophy

**Meta-Ethical Naturalism and the
Fact/Value Dichotomy**

By Fionn O'Donovan

Thesis for the degree of Doctorate in Philosophy

Submitted July 2022

University of Southampton

Abstract

Faculty of Humanities

Department of Philosophy

For the degree of Doctorate in Philosophy

Meta-Ethical Naturalism and the Fact/Value Dichotomy

by

Fionn O'Donovan

Abstract: In this thesis I clarify, and then argue against, the “Fact/Value Dichotomy”, which is, roughly, the notion that there is some sort of fundamental difference between descriptive and evaluative language. To accomplish this, I engage with the arguments of both advocates and critics of this Dichotomy, including Paul Johnston, Hilary Putnam, and Sabina Lovibond. I distinguish between two forms of the Dichotomy and, after identifying weaknesses with extant criticisms of the Dichotomy, give a new argument against it. I then explore the consequences of my argument against the Dichotomy for meta-ethical Naturalism, arguing that my approach allows Naturalists to successfully solve Derek Parfit’s “Normativity Problem”.

Meta-Ethical Naturalism and the Fact/Value Dichotomy: Table of Contents

Thesis Introduction16
Chapter 1, Section 1: The Debate Between MacIntyre and Johnston	20
Ch. 1, S1, part i: MacIntyre on Moral Disagreement and “Emotivism”	20
Ch. 1, S1, part ii: “Emotivism” and Recent Meta-Ethics	23
Ch. 1, S1, part iii: Understanding the Dispute Between the Realist and the Emotivist	27
Ch. 1, S1, part iv: The Role of Johnston’s Account in Defending Emotivism	31
Ch. 1, S1, part v: MacIntyre’s Anti-Emotivist Strategy	34
Ch. 1, S1, part vi: The Johnston-MacIntyre Debate in Relation to this Thesis	37
Chapter 1, Section 2: Johnston’s Wittgensteinian Approach to Moral Philosophy	40
Ch. 1, S2, part i: Johnston’s Methodology	40
Ch. 1, S2, part ii: Johnston’s Embrace of the Fact/Value Dichotomy	43
Ch. 1, S2, part iii: Why Does Johnston Embrace the Fact/Value Dichotomy?	46

Ch. 1, S2, part iv: Johnston's Arguments in Relation to Meta-Ethical Theory	49
Ch. 1, S2, part v: Johnston's Insistence on Ethical Neutrality	54
Ch. 1, S2, Conclusion	57
Chapter 1, Section 3: Can There Be An Ethically Neutral Meta-Ethics?	60
Ch. 1, S3, part i: The Case for the Neutrality Thesis	61
Ch. 1, S3, part ii: A Strategy for Undermining the Neutrality Thesis	66
Ch. 1, S3, part iii: Mulhall's Argument Against the Neutrality Thesis and MacIntyre's Moral/Manipulative Distinction	69
Ch. 1, S3, part iv: Problems for Mulhall's Argument Against the Neutrality Thesis	74
Ch. 1, S3, part v: A New Argument Against the Neutrality Thesis	77
Chapter 1, Section 4: The Ethical Non-Neutrality of Johnston's Wittgensteinian Approach to Moral Philosophy	90
Ch. 1, S4, part i: The Rationale of this Section	90
Ch. 1, S4, part ii: Johnston on Inferences from Descriptive Premises to Evaluative Conclusions	92
Ch. 1, S4, part iii: Classical Teleological Normative Ethical Views	96

Ch. 1, S4, part iv: The Incompatibility of Johnston’s Position and Classical Teleological Views	99
Ch. 1, S4, Conclusion	110
Chapter 2, Section 1: The Fact/Value Dichotomy	114
Ch. 2, S1, part i: The Humean Claim and the Dichotomous Claim: A Clarification	114
Ch. 2, S1, part ii: Thick Concepts and the Humean and Dichotomous Claims	120
Ch. 2, S1, part iii: Separabilism about Thick Language	125
Ch. 2, S1, part iv: Inseparabilism about Thick Language	131
Ch. 2, S1, part v: The Humean Claim, the Dichotomous Claim and Anscombe/MacIntyre Counterexamples	136
Chapter 2, Section 2: Putnam’s Rejection of the Fact/Value Dichotomy	142
Ch. 2, S2, part i: A Summary of Putnam’s Argument	142
Ch. 2, S2, part ii: Problems with Putnam’s Argument	146
Chapter 2, Section 3: Lovibond’s Rejection of the Fact/Value Dichotomy	152

Ch. 2, S3, part i: A Summary of Lovibond’s Argument	152
Ch. 2, S3, part ii: Problems with Lovibond’s Argument	158
Ch. 2, S3, part iii: Desiderata for an Alternative anti-Dichotomy Strategy	166
Ch. 2, S3, Conclusion	168
Conclusion to Chapter 2	169
Chapter 3, Section 1: How to Resist the Fact/Value Dichotomy Without Essentialism	172
Chapter 3, Section 2: Why We Should Reject the Fact/Value Dichotomy	178
Ch. 3, S2, part i: Kovesi on Material and Formal Elements	179
Ch. 3, S2, part ii: Kovesi and the “Purpose Thesis”	183
Ch. 3, S2, part iii: The Implications of the Purpose Thesis for the Fact/Value Dichotomy	190
Ch. 3, S2, part iv: My Argument Against the Dichotomous Claim	197
Ch. 3, S2, part v: What Does My Argument Achieve?	203
Ch. 3, S2, Conclusion:	206

Chapter 3, Section 3: Meta-Ethical Naturalism Without the Dichotomous Claim	208
Ch. 3, S3, part i: Cognitivist Meta-Ethical Theories and the Dichotomous Claim	208
Ch. 3, S3, part ii: Parfit's Objection to Naturalism: The Normativity Problem	211
Ch. 3, S3, part iii: Naturalism Without the Dichotomous Claim	214
Ch. 3, S3, part iv: Differentiating My Naturalism from Jackson's Naturalism	223
Ch. 3, S3, part v: Anscombe/MacIntyre Counterexamples Again	226
Thesis Conclusion: Meta-Ethical Naturalism and the Fact/Value Dichotomy	232
Bibliography	238

Research Thesis: Declaration of Authorship

Print name:	Fionn O'Donovan
-------------	-----------------

Title of thesis:	Meta-Ethical Naturalism and the Fact/Value Dichotomy
------------------	--

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

<p>I confirm that:</p> <ol style="list-style-type: none"> 1. This work was done wholly or mainly while in candidature for a research degree at this University; 2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated; 3. Where I have consulted the published work of others, this is always clearly attributed; 4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work; 5. I have acknowledged all main sources of help; 6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself; 7. Either none of this work has been published before submission, or parts of this work have been published as: [please list references below]: <hr/>

Signature:		Date:	21/12/2020
------------	--	-------	------------

Acknowledgments

I wholeheartedly thank my supervisors, Denis McManus and Philip Stratton-Lake, for their extremely helpful support and advice in writing this thesis. I am also grateful to Rachel Gomez-O'Donovan, Tim Kjeldsen, Ben Paget-Woods, Carla Wiggs, Charlie Boddicker, Felix Hagenstroem, Sigmund Schilpzand, St. John Lambert, Charlotte Unruh, Antonio Scarafone, Genia Schoenbaumsfeld, Alex Gregory, Chris Janaway, Edward Harcourt, and Severin Schroeder for their advice and comments, which have been invaluable. I am grateful to all members of the University of Southampton Philosophy Department between 2016-2021, many of whom have helped with this thesis or with my philosophical development. Finally, heartfelt thanks to my wonderful family, especially Liz O'Donovan, Malcolm Barnsley, Steven Clitherow, John Weafer, Rachel Gomez-O'Donovan, Gaisha Gomez, Basil Gomez, and Joshua Gomez, without whose love and support I could not have completed this project.

Meta-Ethical Naturalism and the Fact/Value Dichotomy

By Fionn O'Donovan

Thesis Introduction

In his 1981 book *After Virtue*, Alasdair MacIntyre attempts to offer an alternative to the widespread “Emotivist” thesis that there can be no rational resolution of moral disagreements. According to “Emotivists”, when it comes to moral disagreements, we are simply stuck in a cycle of repeated assertion and counter-assertion, with no possible way of rationally establishing which among rival moral views is true, or more justified than the others. For MacIntyre, this Emotivist mindset towards morality plays an increasingly dominant role in our life and culture.

In my own attempt to contribute to the philosophical resistance against “Emotivism”, in this thesis I pinpoint, and criticise, one of the major philosophical wellsprings of “Emotivism”: the “Fact/Value Dichotomy” (a term used by e.g. Putnam 2002). I give a precise explanation of what I mean by the “Fact/Value Dichotomy” in Chapter 2, Section 1, but for now, we can sum it up as the idea that there is a fundamental difference in kind between *descriptive* language,¹ the kind of language used to express factual claims like “grass is green” and “smoking causes cancer”, and *evaluative* language, the kind of language used to express, for example, moral claims, like “telling the truth is morally good”. This Dichotomy, as I will explain in more detail below, lends strong support to “Emotivism”. To take a basic example, some meta-ethicists argue that only descriptive

¹ In this thesis, I do not assume any particular position on the issue of the relationship between language and thought. In line with the bulk of the (Wittgenstein-influenced) literature I will be discussing, such as MacIntyre (1981), Lovibond (1983), Johnston (1989) and Putnam (2002), I will orient my discussion primarily around language rather than thought, but do not rule out that my views and conclusions apply also to moral thought, for instance. Whether this will be so will depend on which account of the relationship between language and thought we accept. The reader is encouraged to substitute in their own view on this and decide on the scope of the following arguments (in terms of how they apply to language and/or thought) accordingly.

claims are “truth-apt” (capable of being true or false), and therefore say that evaluative claims, including moral claims,² cannot be true or false. This view then feeds in to the Emotivist thesis, as it tells us that we will not be able to sift truth from falsity, or the more justified from the less justified, in moral debate, as we can in scientific debate, for instance. But if we are not able to do that, it is hard to see how we could rationally resolve moral disagreements.

My principal aim in the thesis is to undermine this “Fact/Value Dichotomy”, and then to explain why excising this Dichotomy from our meta-ethics will improve our understanding of moral language, and the relationship between moral and factual language (see especially Chapter 3, below). By doing so, I hope to thereby contribute to the struggle against “Emotivism”.

Along the way, we will find that, surprisingly, many avowed opponents of “Emotivism” (such as Parfit 2011) also advocate meta-ethical views which take some form of the Fact/Value Dichotomy for granted. Many currently popular versions of Naturalist and Non-Naturalist Realism, I shall argue, fall into this category (including the Naturalism defended by, for instance, Frank Jackson (1998) and the Non-Naturalism defended by Derek Parfit (2011)). Since I reject the Fact/Value Dichotomy, I think opponents of “Emotivism” need to adjust their views, to entirely avoid embracing the Dichotomy. By taking that route, I will argue, we can solve Parfit’s (2011) “Normativity Problem” for meta-ethical Naturalism. By critically engaging with the Fact/Value Dichotomy, then, this thesis produces a concrete, useful result for contemporary meta-ethics.

In Chapter 1, I set the scene by considering the debate between MacIntyre and the arch-Emotivist Paul Johnston (see his 1989), and its relevance for contemporary meta-ethics and issues surrounding moral disagreement (Chapter 1, Section 1). Johnston forcefully defends the Fact/Value Dichotomy, arguing that it captures some crucial conceptual truths about

² In this thesis, I am going to focus on morality and moral claims specifically. I consider moral claims to be one kind of evaluative claim among others (aesthetic, for instance) (cf. Putnam 2002: 19). What I say against the Fact/Value Dichotomy may have implications for non-moral evaluative language and debates too, but I will not be addressing those issues here.

our descriptive and evaluative language (see Chapter 1, Section 2). I argue that even if we do not find Johnston's arguments *prima facie* plausible, it is important for opponents of Emotivism to carefully scrutinise his (unfortunately neglected) arguments. I undermine Johnston's position by arguing that his insistence on the ethical neutrality of meta-ethical theory is untenable (Chapter 1, Section 3), and by arguing that his version of the Fact/Value Dichotomy carries controversial normative ethical commitments which he does not defend (Chapter 1, Section 4).

While these arguments, I believe, show that Johnston's position is unattractive, I then consider, in Chapter 2, Section 1, how a defender of the Fact/Value Dichotomy could hold back from Johnston's radical position while still holding on to the Dichotomy in some form. I distinguish between the stronger form of the Dichotomy defended by Johnston, which I call the "Humean Claim", from the weaker form of the Dichotomy defended by Parfit and others, which I call the "Dichotomous Claim". I consider why the Dichotomous Claim seems more plausible, at least *prima facie*, than the Humean Claim, and how this more moderate position can account for the phenomenon of apparently "thick" (descriptive and evaluative) concepts more easily than the Humean Claim. Hence, the Dichotomous Claim emerges as the most plausible form of the Dichotomy, and the one that critics of the Dichotomy must reckon with, if we are serious about escaping the Dichotomy.

In Sections 2 and 3 of Chapter 2, I consider two previous critical attacks on the Fact/Value Dichotomy from Hilary Putnam (2002) and Sabina Lovibond (1983), respectively. However, I find both criticisms wanting. To sum up briefly, Putnam's argument targets only the Humean Claim, and not the Dichotomous Claim, whereas Lovibond's argument depends on a very strong kind of essentialism about language, which I do not think she adequately defends. Chapter 2 also explores some other criticisms of the Fact/Value Dichotomy and explains why a new argument is needed to combat it. By exploring extant critical accounts, I also generate three desiderata for an improved argument against the Fact/Value Dichotomy, which guide my new argument in Chapter 3.

In the first Section of Chapter 3, I lay out my new strategy for an argument against the Fact/Value Dichotomy, which avoids Lovibond-style essentialism while still being effective against the Dichotomous Claim. Put simply, the strategy involves challenging the advocate of the Dichotomous Claim to make sense of their notion of “thin” (completely non-evaluative) descriptive language, and to motivate scepticism about whether this has been, or indeed can be, done. I then execute this strategy in Chapter 3, Section 2, drawing especially from Julius Kovesi (1967) to undermine the critical plank supporting the Dichotomous Claim: the assumption that the notion of “thin” descriptive language is intelligible. On the contrary, I argue, that notion is better considered a chimera, until or unless the advocate of the Fact/Value Dichotomy gives an adequate explication of it.³

Chapter 3, Section 3 examines the consequences of rejecting the Dichotomous Claim for contemporary meta-ethics. I pinpoint various popular forms of Naturalist and Non-Naturalist realism (supposedly anti-Emotivist views) which assume that the Dichotomous Claim is true. I then consider Parfit’s objection to such Naturalist views, which he calls the “Normativity Problem”: these Naturalist views seem to *explain normativity away*, rather than properly explain or illuminate it. However, I then argue that a version of Naturalism which rejects the Dichotomous Claim does not suffer from this problem. This version of Naturalism also, I argue, explains the linguistic data very well, and does not depend on the dubious Dichotomous Claim. It is hence, I argue, to be preferred to its main rivals.

³ As we will see in more detail in Chapter 3, Sections 1-2, my point is that achieving this may be significantly more difficult than advocates of the Fact/Value Dichotomy give credit for. Moreover, I argue there has not been an adequate explication of the notion of “thin” descriptive language in the literature hitherto.

Chapter 1, Section 1: The Debate Between MacIntyre and Johnston

To give the reader better understanding of the aims and context of my thesis, I will now set out the debate between Alasdair MacIntyre and Paul Johnston. In *After Virtue* (1981), MacIntyre resists the “Emotivist” thesis that it is impossible to rationally resolve moral debates. Johnston later advanced a Wittgensteinian account of moral language which defends “Emotivism” and criticises MacIntyre. For Johnston, descriptive and evaluative language are of “fundamentally different kinds”, while MacIntyre seems to reject that idea entirely. My thesis takes its cue from MacIntyre in this respect, since I argue against “Emotivism” by attacking the Fact/Value Dichotomy which underlies both it and many other popular recent meta-ethical theories.

Part i: MacIntyre on Moral Disagreement and “Emotivism”

In the opening chapters of *After Virtue* (1981), MacIntyre considers the phenomenon of moral disagreement in contemporary society. He discusses examples of arguments on several central moral issues, the sort of arguments which frequently occur both in public and private spheres of our lives: on abortion, war, social justice, and similar topics.

MacIntyre observes that different participants in such arguments tend to appeal to very different concepts (some to “rights”, some to “duties”, some to “virtue”, etc.) and to very different values (some to “justice”, some to “mercy”, some to a perceived need for “strength and success”, some to “equality”, some to “liberty”, etc.) to support their claims about what we morally ought to do (see e.g. MacIntyre 1981: 8-9).

It is unsurprising, MacIntyre says, given the great diversity of ways of thinking about morality and values in our society, and the radically different historical sources of many of

these concepts and values, that we witness frequent, trenchant disagreement on the most burning moral issues (1981: 11).

All this is familiar enough to those with their ear to the ground. More surprising is MacIntyre's claim that contemporary society has "no established way of deciding between these claims" (1981: 9). The result, MacIntyre says, is that our society's moral disagreements "appear to be necessarily interminable" (1981: 9). Since we have no established way of deciding which moral claims are true, or, at least, justified, MacIntyre suggests that, in the long run, our moral arguments invariably peter out into tedious exercises in bare assertion and counter-assertion (1981: 9). With no established way to rationally adjudicate between rival views, where else can moral argument go? If MacIntyre's suspicions are right, it is not just that we *happen* to be unable to reach agreement on these central moral issues. Since (it appears) we lack any means of rationally adjudicating between the views, we lack the ability *in principle* to rationally draw these arguments to a close or reach agreement. Hence why it appears our moral disagreements are *necessarily* interminable.

MacIntyre points out that this situation tends to produce disquiet not just between persons and social groups, but also within ourselves, psychologically (1981: 9-10). Since individuals in our society are typically *aware* of the multiplicity of rival and incompatible moral views on major crucial moral issues, and we appear to lack any rational method for adjudicating between them, one could be forgiven for feeling personally very uncertain about *which* moral beliefs and principles one *should* give one's allegiance to (if any). One understandable consequence would be if individuals become deeply unsure about how to lead their lives, since it is so hard to feel one can rationally settle on a set of moral beliefs and/or attitudes in such circumstances.

One conclusion which can be drawn, if one follows MacIntyre's assessment of our situation, is that, since we cannot settle on any criteria by which to judge the justification of any given moral claim, the decision of whether to accept any given moral claim or principle is, in the end, an arbitrary, non-rational choice. On this view, there is no question

of moral claims being more or less justified, let alone true or false. There are simply various moral views, and it is up to each individual to choose some combination of these by which to live and act. No such choice is more or less rational, or more or less justified, than any other.

MacIntyre (1981) refers to this position as “Emotivism” (see e.g. 1981: 22). The name comes from its being endorsed by certain meta-ethicists, such as Stevenson (1945) and Ayer (1952), who defended a meta-ethical theory known as “emotivism”, on which saying something like “killing is morally wrong” is to express one’s emotional disapproval of killing (and, crucially, does *not* involve stating a truth-apt sentence).⁴ The emotivism of Stevenson and Ayer is one instance of a meta-ethical theory which accepts capital “E” “Emotivism”, the thesis that rational resolution of moral debates is impossible. Terminologically speaking, the reader should note that by “Emotivism” I am not referring to the specific emotivist (small “e”) thesis of Ayer and Stevenson, but rather the broader view about the resolution of moral debate which MacIntyre (1981) calls “Emotivism”.

Emotivism has surprising and bleak implications for moral debate. Participants in moral debate often at least take themselves to be calling attention to genuine *reasons* for accepting certain moral claims, not just expressing their own non-rational attitudes or feelings.⁵ These reasons, at least to the minds of many participants in moral debate, are not just good reasons *because they say, or feel, they are*. Rather, participants in moral debate tend to think that the reasons they identify are good reasons *independently* of what they happen to think or feel about the matter. But if Emotivism is true, then participants in moral debate who take themselves to be identifying reasons of that sort must be mistaken. For if “Emotivism” is true, all that can be going on in moral debate *is* mere non-rational assertion and counter-assertion, or expression and counter-expression, of moral statements. Perhaps some non-rational techniques of persuasion will be added into the mix, and perhaps this will on occasion give the *appearance* that rational argumentation is taking place. But for the Emotivist, it can only amount to just that: an *appearance* of rational

⁴ I discuss this form of emotivism a little further in Chapter 1, Section 3, below, in connection with the possibility of ethical neutrality in meta-ethics.

⁵ See e.g. Smith (1994: 5) for a similar point: “it is a distinctive feature of engaging in moral practice that the participants are concerned to get the answers to moral questions *right*.” Cf. MacIntyre (1981: 10).

argumentation, not the genuine article. Rational resolution of disputes is, on this view, possible in the sphere, say, of the sciences, but not when it comes to moral matters (cf. MacIntyre 2016: 22-23).⁶

In this thesis, I am not going to try to give a knock-down argument against Emotivism, for instance by attempting to prove that the rational resolution of moral debates is possible. Ultimately, the best way of countering Emotivism would be to showcase a method for the rational resolution of moral debates which works effectively. I do not set my sights so high. Instead, I want to achieve two things: first, by undermining some ideas about language and meaning which support Emotivism, and tackling some arguments in favour of it,⁷ I hope to reduce its philosophical appeal. Secondly, while I do not supply a method for the rational resolution of moral debates, I do want to add plausibility to the idea that rational resolution of such debates is *possible*. For instance, the Naturalist meta-ethical position I set out in Chapter 3, Section 3, below, suggests that the Emotivist idea that rational resolution of debate is possible in the sciences, but not in the area of morality, is mistaken. If rational resolution of debate is possible in the sciences – as it seems safe to assume it is – then this thesis suggests that it should, in principle, be possible in the moral sphere too.

Part ii: “Emotivism” and Recent Meta-Ethics

As MacIntyre (1981: 22-23) points out, the Emotivist view is, for many of us, difficult to believe (cf. Streumer 2017). Many meta-ethicists believe it is more plausible to take moral debate at face value, accepting that those participating in it are making claims that could be true or false, and rationally more or less justified (as they usually take themselves to be), and that they are trying (successfully or otherwise) to rationally defend these claims (see e.g. Smith 1994 for further discussion of this point).

⁶ I return to, and reject, this claim later (see e.g. Conclusion).

⁷ See, for example, my criticisms of Johnston’s (1989) explicitly “Emotivist” position in Chapter 1, below, and my criticisms of Parfit’s view on descriptive and evaluative language in Chapter 3, Section 2. While Parfit is not an “Emotivist”, I will argue that his views on descriptive and evaluative language concede too much to the “Emotivist”.

The upshot of this is that many meta-ethicists cling on to realist views, which are apparently inconsistent with Emotivism, and which conform with the idea that, when making moral claims, we are making claims that are capable of being true or false and, at least in principle, possible to rationally justify (i.e. to show to be true, or likely to be so). So, for instance, on Moore's (1903) realist view, the claim "telling the truth is morally good" is true just in case acts of truth-telling possess the non-natural moral property of goodness, in a way similar to how "grass is green" is true just in case grass possesses the natural property of greenness.⁸

While such realist views remain popular, there is still significant popularity and appeal to meta-ethical views which, unlike Moorean realism, conform with Emotivism.

An influential and important example of such a view is Allan Gibbard's (1990) non-cognitivist expressivism. According to Gibbard, when we state some moral claim, like "telling the truth is morally good", we are expressing our state of mind that truth-telling is desirable. "Good things are desirable, and the better of two things is the one that is preferable", he writes (1998: 241). Gibbard cashes out the notion of "preferability" in terms of "rationality": "the preferable thing is the one that it would be rational to prefer" (1998: 243).

On the face of it, then, Gibbard is, unlike the Emotivist, happy to allow talk about the "rationality" of different moral claims. Telling the truth may be "preferable" to telling lies, and, if so, it would be "rational" to prefer truth-telling to telling lies, on Gibbard's view. Should we say, then, that Gibbard denies the Emotivist thesis that there can be no rational adjudication between moral claims, and hence denies that the choice of whether to endorse a given moral claim is, in the end, a matter of non-rational arbitrary personal choice?

⁸ I discuss Moore's view in more detail in Chapter 1, Section 2, below.

We should not be too quick to distance Gibbard's expressivism from Emotivism, because, crucially, for Gibbard the notion of "rationality" is itself subjected to the expressivist analysis: "someone who calls something "rational" is therefore expressing his state of mind" (Gibbard 2009: 161-2). He also writes that

We explain the term [rational] by saying what state of mind it expresses. In this sense the analysis is expressivistic ... The analysis is non-cognitivist in the narrow sense that, according to it, to call something rational is not to state a matter of fact, either truly or falsely (Gibbard 1990: 8).

This latter quotation makes clear that, for Gibbard, what it is to call something "rational" is to be cashed out in strictly non-realist terms. On his account, calling something "rational" is to express one's state of mind towards it, and *not* to state a matter of fact.

So while Gibbard allows that there is a legitimate place for talking, for example, about it being rational to prefer truth-telling to lying, in a way he only pushes the issue at stake between the realist and the "Emotivist" back one step, since, on his view, saying that "it is rational to prefer truth-telling to lying" is *itself* expressing one's state of mind, i.e. that one is more favourably inclined towards truth-telling than lying. The result is that, on Gibbard's view, if two persons, in a debate about truth-telling, ended up conflicting, with one person saying "it is rational to prefer truth-telling to lying", and the other sincerely denying this, there would be no way of rationally settling this disagreement. We would, rather, simply have a fundamental difference in attitudes, which would not be amenable to rational resolution.

As MacIntyre (2016: 21) comments, this means that, in Gibbard's view, there are very stringent limits to the role "rationality" plays in the best account of moral disagreement. This aspect of Gibbard's view places him, I believe, firmly in the Emotivist camp. For consider, by contrast to Gibbard's view, a realist view on which there *is* a fact of the matter about the goodness of truth-telling. For such a realist, a disagreement about whether truth-telling is morally good is *not* simply a matter of two or more persons having different states of mind or attitudes towards truth-telling. It is rather a matter of two or more persons

disagreeing about the facts, in much the same way as a disagreement about whether grass is green is a disagreement about the facts (MacIntyre 2016: 24-5). Moreover, to say that “it is rational to prefer truth-telling to lying”, on such a realist view, is not simply to express one’s state of mind towards truth-telling, but also to convey the thought that it is *true* that truth-telling is better than lying and that there is a justification for believing this to be so – a justification that justifies one’s approval of truth-telling, rather than just re-asserting one’s approval of truth-telling.

The reason I am focusing on this difference between Gibbard’s expressivism and (typical)⁹ realist views is to convince the reader that, if we scratch the surface of Gibbard’s expressivism, we find that it is Emotivist (cf. MacIntyre 2016: 22). Gibbard allows that talk of the “rationality” of endorsing a moral claim is meaningful, but only in a *much* thinner sense than the realist wants. It does not take long for the disagreement between the realist and the Gibbardian expressivist to return to the disagreement MacIntyre (1981) calls to attention between the Emotivist who thinks that, ultimately, moral disagreements are *not* amenable to rational adjudication, and the realist who thinks that they *are* so amenable, so to speak, all the way down.

While various kinds of realist views remain popular, so do Emotivist views. What accounts for the enduring persistence of meta-ethical views inspired by Emotivism? Partly, I think it is that while realism gives a seemingly more accurate picture of moral language and moral debate than Emotivists,¹⁰ Emotivists nevertheless appear to be right that we lack an established procedure for rationally adjudicating moral disagreements.

Hence, while the realist may give a more accurate picture of moral language, the realist seems too optimistic about the possibility of rationally resolving moral debates. It is unsurprising, then, that views such as Gibbard’s, which try to capture realist insights about

⁹ The reason I say “typical” realist views is that some meta-ethical views which seem to count as realist may still be consistent with “Emotivism”. See e.g. MacIntyre’s discussion of Hare’s views in MacIntyre (1981: 23-24), and cf. Eklund’s discussion of “ardent realism” in his (2017: 1-2). As we will see in Chapter 3, Section 3, many forms of realism take the Fact/Value Dichotomy for granted, which I will argue is a mistake and concedes too much to the “Emotivist”.

¹⁰ Cf. Smith (1994).

moral language while avoiding the realist ontology of moral truths and moral facts, have widespread appeal (cf. Blackburn's "quasi-realism" (1984, 1998) as well as Gibbard's work (e.g. Gibbard 1990, 2003); for an overview of the gradual compromises and accommodations between realists and expressivists see e.g. Hare 2001: Ch. 1 and Dreier 2004).

In this thesis, I aim to strengthen the realist's position in this struggle with the Emotivist. Pessimism about the possibility of rationally resolving moral debates is, I argue, motivated in large part by a widespread but erroneous view about moral language and "non-moral" descriptive language, and how they relate, a view which I refer to with Hilary Putnam's (2002) name "the Fact/Value Dichotomy". I will argue, however, that many realists (e.g. Parfit 2011) accept this Dichotomy, thereby conceding too much to Emotivism. By undermining this Dichotomy, we can show why the Emotivist's pessimism about the possibility of rationally resolving moral debates is unfounded, and why realists should not concede as much to Emotivism as they have often elected to.

Part iii: Understanding the Dispute Between the Realist and the Emotivist

While expressivists like Gibbard, then, make a small concession to the realist by allowing talk of the "rationality" of accepting a moral claim, this does little to resolve the deeper underlying disagreement between the realist and the Emotivist. Ultimately, expressivists continue to defend the Emotivist idea that whether or not one accepts a moral claim is simply a matter of one's "fundamental cast of mind" (MacIntyre 2016: 21). It is not, on the expressivist view, as though one can "dig deeper" and question whether those fundamental casts of mind are *themselves* rational, or justified, for to call something "rational" or "justified" is, according to expressivism, no more and no less than to express one's state of mind towards it (MacIntyre 2016: 21, cf. Blackburn 1998: 241). This will, of course, be anathema to most realists. The realist and expressivist will, then, continue to disagree on whether it is possible to rationally resolve moral debates. The realist will not be satisfied with Gibbard's view that a dispute over whether "it is rational to prefer truth-telling to

lying” (for example) amounts to the mere expression of one state of mind, on which truth-telling is preferable to lying, and the counter-expression of another incompatible state of mind, on which it is not.

A crucial part of the dispute between the realist and the Emotivist, then, is whether it *is* possible to secure some method of rationally resolving even the deeper moral disagreements such as the example we discussed about the rationality of judging truth-telling to be morally better than lying. It is important to pay attention to the word “possible” here. To vindicate realism and move away from Emotivism, we do not necessarily have to find a specific method which reliably resolves moral debates rationally (although that would, of course, be the best-case scenario for the realist). The realist can make do with defending the claim that, contrary to Emotivism, the rational resolution of moral debates is *possible*. This would undermine the Emotivist thesis and remove the initial motivation behind non-cognitivist expressivist accounts. The reason for this, I have suggested, is that expressivism and other Emotivist-inspired views take part of their motivation from pessimism about the possibility of rationally resolving moral disagreements.

To make clearer the difference between realism and Emotivism, consider MacIntyre (2016) on the difference between how a person who accepts Emotivism would and could reflect on their desires, and how a person who accepts an Aristotelian form of realism would and could do so. MacIntyre imagines a young person reasoning about an internal moral conflict they are having: on one hand, they believe they have good reasons to tell an uncomfortable truth to their family; but on the other hand, they have a very strong desire *not* to do so (my example here is adapted from MacIntyre’s discussion at 2016: 32-4).

If this young person is an Emotivist (or expressivist), “it would at once occur to her that her conflict is between one part of herself and another” (MacIntyre 2016: 32), and, we might add, *simply* between one part of herself and another (i.e. *not* between her desires and the truth about what she ought to do, for instance). Her judgment that she has good reasons to tell the truth here is itself a manifestation of some part of her overall cast of mind which

approves of telling the truth in this case. Her strong desire *not* to tell the truth in this situation, plainly enough, issues from some other part of her overall cast of mind which *disapproves* of telling the truth in such circumstances. What reason might she have, though, for siding with one part of her psyche rather than the other?

On the Emotivist picture, the basic idea will be that she should settle on whichever action she ultimately has a stronger desire to perform. According to the expressivist, the only meaningful sense in which she can come to judge it more rational to either tell the truth or not is by coming to realise that her overall cast of mind favours one course more than the other. Whichever action she ultimately has a stronger desire to do is the one which, from the expressivist point of view, she morally ought to do.

If this young person was thinking in Aristotelian realist terms, MacIntyre says, the picture would be quite different. In this case, it would be possible for the young person to judge that her strong desire not to tell the truth is a *misdirected* desire (2016: 34). Desires, the thought goes, aim at some good, or goods. According to Aristotelian realism, moreover, not all goods are created equal: some are more valuable and more worthy of pursuit by rational human beings than others.¹¹ And in saying this, the Aristotelian does *not* simply mean to express her state of mind that these goods are more desirable, but rather to make a true claim about the nature of the good in question. So for example, perhaps one's health is a more valuable good than the good of the pleasure one may get from eating a large chocolate cake. Nevertheless, it is possible that one's state of mind is overall such that one desires to eat the cake more strongly than one desires to stick to one's diet for the sake of one's health. For the Aristotelian, this would be a case of *misdirected* desire: one's desires are not, morally, as they should be, since the virtuous agent would have a stronger desire for the more worthy good (in this case, health). So on the Aristotelian realist view, unlike on the expressivist view, one's desires and other states of mind are *accountable* to an independent standard of goodness (cf. Eklund 2017: 1 on "ardent realism").

¹¹ Cf. my discussion of "Classical Teleological" ethics in Chapter 1, Section 4.

However, this realist idea that goods can be evaluated by standards which are independent from our desires and other states of mind is exactly what the Emotivist, and the expressivist, call into question. How can we adjudicate whether health is a more worthy good than that of the pleasure of eating sweet food? What else can we appeal to, the Emotivist asks, to settle such a matter, other than our own desires and other relevant states of mind of ours? The Emotivist rejects the realist's idea that there is, *or could be*, a way of rationally resolving such disputes, beyond the kind of appeal to desires or states of mind that Gibbard allows for.

Hence, if we *cannot* somehow vindicate the possibility of rational resolution of moral debates, this will add serious weight to the Emotivist view that, although it may make sense to talk about the "rationality" of endorsing a moral claim, this can involve no more than expressing one's state of mind towards that claim. So moral disagreements about what it is rational to approve of amount to no more than the expression and counter-expression of different states of mind. If this is right, then the realist's talk about moral truth, and the rational adjudication of such deeper moral disagreements, is just that: talk. This is, I believe it would be fair to say, what the Emotivist wanted to conclude all along (cf. MacIntyre 2016: 23-26).

Moreover, if we cannot secure such a method, the apparent plausibility of realism as an account of moral language will increasingly be undermined by Emotivist accounts such as Gibbard's, especially because expressivist views like Gibbard's look increasingly easy to reconcile with common-sense views about moral language, at least compared with Stevensonian emotivism.¹² (The broader story about expressivist attempts to make their account compatible with ordinary language involves, among other things, their responses to the Frege-Geach problem, on which see Schroeder 2008 and MacIntyre 2016: Ch. 1). In this case, realism will remain vulnerable to its Emotivist critics. The arguments of this thesis aim to safeguard realism from this Emotivist criticism and help push back against Emotivism.

¹² See Stevenson (1945). I discuss Stevenson's emotivism in more detail in Chapter 1, Section 3, below.

Part iv: The Role of Johnston's Account in Defending Emotivism

What would make Emotivism more plausible is an account of moral language, or an argument, which gives reasons to think that rational adjudication of rival moral views is, in principle, impossible. While many meta-ethicists think that the mere existence of persistent, widespread disagreement is good evidence for Emotivism,¹³ if there is hope that we (realists) have simply not *yet found* the right method for assessing moral truth, we could reasonably hold out hope that there *are* moral truths, and it is just that we need to find the right method for determining what they are. This important point in defence of realism is stressed by David Enoch (2011: 35).¹⁴

Enoch's point is especially apposite because, even though expressivists like Gibbard and Blackburn have made Emotivist views more compatible with ordinary language, there is still a widespread sense that realist views are *more* easily compatible with ordinary language. Emotivists are hence often regarded (rightly or wrongly) as being at a slight initial dialectical disadvantage compared to the realist, and are left needing to compensate for this in other areas. As mentioned in the previous Section, one common aim such theorists have is to give an account which does not rely on the seemingly heavy ontological commitments that realism entails (see e.g. Ayer 1952: 106-7, Blackburn 1998: 296). The realist's idea that moral debates can be rationally resolved seems to require the existence of things such as moral facts or moral properties, about which Emotivists are sceptical (see e.g. Mackie 1977).

Even if we sympathise to some degree with this scepticism, however, we should ask what, if anything, licenses the Emotivist's (seemingly extreme) conclusion that we fail to rationally resolve our moral disagreements because such disagreements are *necessarily* impossible to rationally resolve. Might it not be, in the spirit of what Enoch (2011)

¹³ See e.g. Mackie (1977: 36), and for a survey of the literature on this point, see Gowans (2000a).

¹⁴ I return to Enoch's point in relation to the debate between realists and "Emotivists" in my Conclusion, below.

suggests, that we simply have yet to find the right method for detecting moral truths, which has (often if not always) frustrated our efforts so far, but need not do so in future?

As a result, I think Emotivists, including expressivists, need an account that makes their denial of the possibility of rational adjudication between rival moral views plausible. It is not that expressivism, for example, is taken to *prove* that rational adjudication between rival moral views is impossible. It is the other way: the thought that such a thing is impossible is, rather, something that motivates expressivism. The plausibility of Emotivism thus hinges on the plausibility of the denial that rationally resolving moral debates is possible. Equally, this leads to a way of opposing Emotivism: if we can undermine the attempt to prove the impossibility of rationally resolving moral debates, this will be a significant setback for Emotivists (see, in this regard, my arguments against Johnston in Chapter 1 and my argument in Chapter 3, Section 2, below).

It is in offering such a principled justification of Emotivism that Paul Johnston's (1989) account of moral language has much of its (underappreciated) importance. Johnston, commenting on MacIntyre's diagnosis of the apparent state of disorder in moral discussions in our culture, argues that these are necessary features of moral argument, in *any* cultural context. For Johnston, as we shall see in the next Section, the rational resolution of moral debates, and the philosophical justification of a moral claim, are both impossible *in principle*. Hence Johnston, by MacIntyre's broad definition (the same one I am using in this thesis), is an Emotivist. What differentiates Johnston from other Emotivists is that he reinforces his Emotivism by explicitly arguing that descriptive language and evaluative language are, semantically, of two fundamentally different kinds. He argues that this conceptual truth undermines the very possibility of giving evidence for, or proving, or rationally justifying (in the heavy-duty sense the typical realist wants¹⁵), a moral claim.¹⁶

¹⁵ i.e. a kind of justification that goes beyond Gibbard's "rational justification" by appeal to one's own state of mind.

¹⁶ I explore Johnston's arguments in detail in Chapter 1, Section 2.

Since Johnston is an Emotivist, he rejects MacIntyre's suggestion that the level of conflict in, and the apparent interminability of, moral argument in our culture is cause for concern (Johnston 1989: 88). Johnston instead calmly assures us that these patterns are exactly what one would expect, if one had a clear understanding of the nature of moral language (i.e. an Emotivist understanding) (see 1989: 89). For Johnston, as I will explain in more detail in the next Section, evaluative language (including moral language) is different in kind to descriptive language (the kind of language we use to state ordinary empirical facts like "grass is green"). In Chapter 2, Section 1, I will show that Johnston's view hence falls under the umbrella of the "Fact/Value Dichotomy". The special features of descriptive language, Johnston argues, mean that rational resolution of disagreements is possible when the claims at stake are descriptive, but the special features of evaluative language (he maintains) equally guarantee the impossibility of such resolution in the case of evaluative claims.

Since Johnston's account seems to supply the principled justification of Emotivism which Emotivist meta-ethical theories need, the debate between MacIntyre and Johnston warrants further investigation. If the reader is wondering why I am choosing to focus on these philosophers, rather than more recent authors and their (perhaps more refined) meta-ethical stances, it is because I believe tackling the Johnston-MacIntyre debate will help us expose erroneous assumptions (particularly those which form part of the Fact/Value Dichotomy) which have led many more recent meta-ethicists in the wrong direction (see Chapter 3, Section 3 for more on this point).

What Johnston fails to engage with, in his critical response to MacIntyre, is MacIntyre's own alternative explanation of our culture's apparent inability to rationally adjudicate between the many rival moral perspectives present within our culture, which is incompatible with Johnston's. I will now sketch MacIntyre's alternative account.

Part v: MacIntyre's Anti-Emotivist Strategy

According to MacIntyre, our culture's inability to rationally adjudicate rival moral views is *not* a result of the essential impossibility of rationally justifying any moral claim (as Emotivists like Johnston suggest), but rather a historically and culturally specific limitation resulting from contingent historical and conceptual circumstances (see e.g. MacIntyre 1981: x, and 136-137). For MacIntyre, that is, rationally justifying moral claims may *become* impossible in certain times and places, but it is by no means impossible in *all* times and places.

Crucial, of course, for MacIntyre's rival explanation to be at all plausible, is for him to give a history of our current situation which explains why moral arguments in our culture have appeared to become rationally interminable. But even more crucial (especially for answering a critically-minded Emotivist like Johnston) is for MacIntyre to give examples of cultures in which rational resolution of moral arguments was, or is, possible. For it is exactly this possibility which Emotivists deny.

In order to deny Emotivism, MacIntyre tries to answer both the challenges mentioned in the above paragraph (see MacIntyre 1981, 1988). Although it is impossible to consider MacIntyre's efforts in full detail here, for our purposes it is important to consider the case of Homeric Ancient Greece as an example of a culture in which, for MacIntyre, rational resolution of moral debates was possible.

According to MacIntyre, Homeric Ancient Greece is an example of a culture in which a host of concepts crucial to morality were commonly understood to work in a *functional* way. So, for example, MacIntyre tells us how in the Homeric system, certain social duties, e.g. providing economic security for family, was considered part of the *function* of a "man". The "just" or "virtuous" man is hence by definition one who could carry out such social functions effectively, and in this cultural context to understand "justice" just is to

understand that just actions are those which effectively lead to these functions being carried out.

As MacIntyre points out (1981: 214-5), the Homeric Greek example is one in which moral concepts have a role in a teleological system. A given “end” (or “*telos*”) – say, for human beings – provides a standard of evaluation for whether some act was “right”, “just”, etc. For within this system, whether some act is right depends on whether it promotes this end, or *telos*.

The crucial aspect of this teleological structure of Homeric moral thought, for our purposes, is that it involves a highly different way of thinking and speaking, not only about moral truths, but also about the relation between factual and moral claims, compared to our (Emotivist-influenced) modern culture. So, within this teleological structure, MacIntyre says (1981: 69) it makes perfect sense to speak of moral claims being true or false. There can be, for example, known and knowable facts about which actions and character traits do as a matter of fact tend to help realise certain ends, and which do not. Thus, within this system, there are *facts* about which actions and character traits are “right”: they are just those which lead (most reliably) to the *telos*.

Let me elaborate on why, for MacIntyre, thinking about concepts such as “man” functionally leads to a radically different way of understanding the relationship between factual and moral claims, compared with the Emotivist’s. For Johnston, owing to the logical dichotomy between facts and evaluations, there can be no logical deductions from factual descriptive claims to evaluative moral claims. There is simply a “logical gulf” (to use Meynell’s (1971) phrase) here which cannot be bridged, Johnston thinks. MacIntyre, by contrast, denies this.

For MacIntyre, a deduction from factual descriptive claims to evaluative moral claims *is* perfectly possible, so long as at least one of the concepts appearing in the premises is *functional* (1981: 69). So, e.g., MacIntyre says, from the descriptive premise “the hands on

your watch no longer move” we can deduce the evaluative conclusion “your watch is a bad watch”, and from the factual descriptive claim “farmer X attains the highest crop yields per square metre in his region” we can deduce the evaluative conclusion “he is a good farmer” (MacIntyre 1981: 68). The point here is that the concept “watch” appears to be functional in the sense that, if one understands what a watch is, one must understand that a good watch tells time accurately, and to understand the concept “farmer”, one must understand that a good farmer produces high crop yields efficiently (cf. Julius Kovesi’s (1967) account of “formal elements” in Chapter 3, Section 2, below).

For MacIntyre, such examples reveal that there is no “unbridgeable logical gulf” between factual descriptive claims and evaluative moral claims, contrary to Johnston and other Emotivists. The ease with which it is possible to pass between descriptive factual claims and evaluative moral claims, MacIntyre claims, shows that it is only in specific cultural contexts that persons begin to cease to see any relation between factual descriptive claims and evaluative moral claims, and accordingly to grow more sceptical that there is any way of rationally differentiating true from false moral claims: specifically, those contexts in which the crucial concepts deployed in moral arguments are no longer considered to be *functional* concepts.

If MacIntyre is right, there would be significant consequences for the rational resolution of moral disagreements. For example, if factual premises really can entail moral conclusions, then it seems rational agreement on at least some moral questions should be equally as achievable as rational agreement on scientific matters – contrary to Emotivism (MacIntyre 2016: 24-25). Whereas the Emotivist draws a sharp line between empirical, scientific debate and moral debate, then, MacIntyre’s view challenges the division between empirical, rationally resolvable debate and moral debate, which is supposedly not rationally resolvable. The present thesis supports MacIntyre’s position against Emotivism. The key, however, is to see how and why MacIntyre’s insights support a form of meta-ethical Naturalism which *entirely avoids* the Fact/Value Dichotomy, which many standard contemporary forms of Naturalism problematically fail to do (see my argument in Chapter 3, Section 3, below).

Part vi: The Johnston-MacIntyre Debate in Relation to this Thesis

On the basis of his counter-examples to the “no-evaluative-conclusions-from-descriptive premises” principle, and his observation that certain cultures understood moral concepts functionally, MacIntyre rejects Johnston’s conclusion that descriptive claims and evaluative claims are of fundamentally different logical kinds. For MacIntyre, evaluative moral claims “just are a kind of factual statement” (1981: 173). As I just explained, this controversy about the relation between factual and descriptive language is crucial to settling the dispute between MacIntyre, and all others who wish to vindicate the possibility of rational adjudication between moral claims, and Emotivists, who deny this possibility.

The reason I sketched MacIntyre’s anti-Emotivist strategy above is that I believe it helps identify a way of pushing back against Emotivism that I can develop further in this thesis. Clearly, MacIntyre and Johnston have diametrically opposed accounts of descriptive language, evaluative language, and the relationship between them. By undermining Johnston’s account of these matters and buttressing MacIntyre’s, I will support MacIntyre’s argument that the Emotivist mistakes our *contingent* inability to rationally resolve moral disputes with a necessary inability to do so, thereby vindicating, at least to a significant degree, Enoch’s optimism that some method of rationally resolving moral debates is possible.

Hence, in the remainder of Chapter 1, I tackle Johnston’s arguments in favour of Emotivism and the Fact/Value Dichotomy, with the aim of exposing problems with his position: first, that he cannot deliver on the ethical neutrality he promises (see Section 3, below) and second, that his position commits him to normative ethical claims that he neither acknowledges nor argues for (see Section 4, below).

While, I argue, these are serious problems with Johnston’s position, they do not in themselves undermine the account of descriptive and evaluative language which really

drive Johnston's Emotivism. From Chapter 2 onwards, I tackle the underlying view of descriptive and evaluative language which drives Emotivism: the Fact/Value Dichotomy.

While MacIntyre's "counterexamples" to the Fact/Value Dichotomy, namely the examples of the watch and farmer mentioned in the previous part of this Section, pose an important challenge to Emotivists, I will show that they are not sufficient on their own to undermine the Fact/Value Dichotomy. Crucially, advocates of the Fact/Value Dichotomy have resources with which to answer MacIntyre's cases, including a "Separabilist" strategy which involves trying to break concepts like "watch" and "farmer" into separate descriptive and evaluative components (see Chapter 1, Section 4 and Chapter 2, Section 1, below). The tenability of the Dichotomy hence depends on whether these responses to MacIntyre's cases can be sustained. By developing a new argument against the Dichotomy, I shall pinpoint why these defences are not sustainable, thereby undermining Emotivism.

Moreover, I aim to show that my new way of undermining the Fact/Value Dichotomy has important consequences for contemporary meta-ethics. Once we realise that the idea of a "pure" (evaluation-free) descriptive language is a chimera, we clear the way for a more plausible meta-ethical Naturalist position. I motivate this position by comparing and contrasting it with the views of Parfit (2011), and using it to solve Parfit's "Normativity Problem" for meta-ethical Naturalism (see Chapter 3, Section 3).

I argue that many of the popular recent views in recent meta-ethics, such as Jackson's (1998) Naturalism and Parfit's (2011) Non-Naturalist realism, either explicitly or implicitly depend on at least some Johnstonian ideas about the logical differences between factual descriptive claims and evaluative moral claims.¹⁷ Investigating and developing MacIntyre's radically different approach will help me argue that these underlying Johnstonian ideas should be excised from our meta-ethical theories, and much of the motivation behind "Emotivism" will thereby be undermined. Moreover, refusing to concede ground to Johnston will allow us to take up a form of Naturalism which, unlike

¹⁷ See Chapter 3, Section 3, below.

conventional Naturalist views, does not suffer at the hands of Parfit's "Normativity Problem".

Moreover, this thesis will defend MacIntyre's claim that evaluative moral claims "just are a kind of factual statement" (1981: 173), setting out more precisely the sense in which this is true. Excising the Fact/Value Dichotomy, I will argue in Chapter 3, allows us to avoid the flaws which Parfit rightly identifies with *reductivist* Naturalism (e.g. Jackson 1998) without giving up on Naturalism altogether.

I will now turn to my first task, which is to examine Johnston's arguments in defence of Emotivism in more detail.

Chapter 1, Section 2: Johnston's Wittgensteinian Approach to Moral Philosophy

The purpose of this Section is to summarise and explain Johnston's Wittgensteinian¹⁸ approach to moral philosophy and his reasons for embracing the "Fact/Value Dichotomy", as set out in his 1989 book *Wittgenstein and Moral Philosophy*. I begin by explaining Johnston's Wittgensteinian methodology (part i). I proceed to examine his main arguments that descriptive statements and evaluative statements are "of fundamentally different kinds" (part ii)¹⁹ and that so-called "second-order" enquiry in moral philosophy can be conducted in an entirely ethically neutral way (part iii).²⁰

As discussed in the previous Section, Johnston's defence of the Fact/Value Dichotomy is important because it props up the Emotivist view outlined, and criticised, by MacIntyre (1981). I will be arguing that finding the right way of rebutting Johnston's arguments opens up an improved understanding of the relationship between factual descriptive language and evaluative moral language, and thereby allows us to move beyond Emotivism.²¹ First, though, we need to understand Johnston's reasons for accepting the Fact/Value Dichotomy.

Part i: Johnston's Methodology

To explain the methodology of his Wittgensteinian moral philosophy, Johnston introduces a distinction between "descriptive" and "revisionary" philosophy (1989: 10-11).²² Johnston explains the notion of descriptive philosophy by referring to Parfit, who explains that such philosophy describes and explains the "unchanging central core in our beliefs about

¹⁸ I would dispute whether Johnston's position is genuinely Wittgensteinian. However, I will not pursue this exegetical matter further in this thesis.

¹⁹ I argue against this view in Ch. 3.

²⁰ I argue against this view later in this Chapter: see Ch. 1, Section 3, below.

²¹ See Ch. 3, Section 3, below.

²² The distinction comes from Strawson (1959: 9). To clarify, none of my own arguments will appeal to or depend upon this distinction: I mention it because Johnston uses it to explain his view.

ourselves, and the world we inhabit” (Parfit 1984: x). It is aptly characterised as “descriptive” because it aims at accurately elucidating the “central core” of our beliefs. Revisionary philosophy goes beyond this, because it does not rest content with simply setting out and accurately understanding the “central core” of our beliefs: it also subjects them to critical scrutiny, and aims at changing and improving our core beliefs (where it can find good reasons for doing so). The crucial difference between descriptive and revisionary philosophy, Johnston explains, is that only the latter is *evaluative*, since, unlike descriptive philosophy, it aims at “improving” the “central core” of our beliefs in some way.

Johnston states that his approach will be descriptive, not revisionary (1989: 11). However, he also states that he provides a specific *kind* of description, namely “the clarification and systematic representation of the rules which define our concepts” (1989: 11). Johnston, following Wittgenstein (1953: Section 90), calls this “grammatical” investigation, and sometimes also uses the words “logic” or “logical”, in a similar vein. To avoid exegetical debates about the notions of “grammar” and “logic” in Wittgenstein, I will not follow Johnston in deploying those terms, but I will, like Johnston, distinguish between empirical description, on the one hand, and description and clarification of our concepts, on the other.

Empirical description, Johnston says, “will typically involve the weighing-up of evidence” (1989: 11) to ascertain whether descriptions are true or false. “In contrast”, Johnston says, “... a grammatical investigation involves neither the assembling of evidence nor its assessment. Rather, this type of investigation is an exercise in clarification” (1989: 11).

Johnston believes that, by clarifying our moral language and concepts, confusions and difficulties in moral philosophy can be dissolved (1989: 200). For instance, Johnston thinks that the philosophical quest to find an “objective proof” of some of our moral beliefs is the product of conceptual confusion (1989: Ch. 5). Once we realise the grammatical difference between descriptive and evaluative statements,²³ Johnston says, we will realise that the notion of “objectivity” is at home only in the field of factual, empirical language, and give

²³ I discuss this component of Johnston’s position in part ii) of this Section.

up the confused idea that “objectivity” could be sensibly applied in moral contexts. (Note the clear affinity here with the Emotivist position discussed in the previous Section).

After dissolving confusions in this way, Johnston thinks, we end up with a clearer understanding of our own lives and moral language, and are hence less prone to misconceiving this reality due to the misleading influence of philosophical “pseudo-questions” (1989: 200).²⁴

Johnston rules out trying to “advance empirical hypotheses about the origin of ethics” (1989: 17), as he thinks philosophers like Russell (1918) do.²⁵ This is why the kind of description Johnston is interested in is not empirical but “grammatical”: it involves description and elucidation of our concepts, specifically our moral concepts. As we will see in more detail in part iii) of this Section, Johnston also sharply distinguishes his descriptive investigation from any sort of revisionary philosophical project. Johnston describes his approach as “non-substantive in the sense [that it involves] simply describing our concepts and practices without seeking to justify them” (1989: 25). He maintains that his approach requires a kind of “self-discipline” to “restrict oneself to ‘mere’ description” (1989: 25), that is, to avoid the temptation to try to justify the concepts and practices we describe.

So, in summary, Johnston’s methodology is descriptive in the sense that he aims to clarify our moral concepts: to describe our moral language and concepts and thereby dissolve philosophical confusions about morality. It is also, in his view, non-revisionary and non-evaluative, since it aims only at conceptual clarity and not at any sort of justification or vindication of our moral concepts, practices or beliefs.²⁶

²⁵ See Johnston (1989: 16-17) and Russell (1918: 108), who claims “ethics is in origin the art of recommending to others the sacrifices required for co-operation with oneself”. As Johnston points out, this is simply an unsupported, speculative empirical claim.

²⁶ Whether the ethically neutral path Johnston tries to follow is viable will be questioned in Ch.1 Section 3, below.

Part ii: Johnston's Embrace of the Fact/Value Dichotomy

It is worth re-emphasising that the central thesis of Johnston's book is that there is an irreducible conceptual distinction between descriptive claims and evaluative claims²⁷ (see e.g. 1989: 201). (For Johnston, as for the present author, moral claims are a sub-set of evaluative claims). Johnston claims that "statements such as "this is good" and "this is red" are of fundamentally different kinds" (1989: 95-6). The idea here is that "this is good" is a paradigmatic evaluative claim and "this is red" is a paradigmatic descriptive claim,²⁸ so Johnston is here expressing his core thesis that descriptive and evaluative claims are of fundamentally different kinds.²⁹ He further tells us that "moral judgments... belong to a different category from empirical judgments and... have a completely different grammar" (1989: 201). On his view, "the contrast between fact and value... is thus revealed to be grammatical in nature" (1989: 84). Johnston makes this point to clarify that his thesis about descriptive and evaluative language is intended to be strictly a *conceptual* one, rather than one that commits him to any metaphysical claims.³⁰

Since Johnston argues that factual descriptive claims "have a completely different grammar" to evaluative moral claims, it seems safe to say that Johnston accepts the "Fact/Value Dichotomy."³¹ What exactly is it, though, to accept the "Fact/Value Dichotomy"? Since this is such a crucial question for my thesis, I will delve into this issue in more detail in Chapter 2, Section 1. For now, it will suffice to note that I am using the

²⁷ Johnston sometimes uses the word "statements" instead of "claims". I have decided to use the word "claims", while some meta-ethics literature speaks of "normative judgments" (see e.g. Gregory 2017). I do not consider there to be a significant difference between these terms, since the main reason such terms are used is to avoid calling sentences like "this is good" *propositions*, which would controversially assume that they can have truth-values (an idea disputed by many expressivists, see Ch. 1, Section 1, above). In other words, the main point is that we can agree that we need some word to refer to the kind of sentences that meta-ethicists wish to analyse, without begging the question against any serious meta-ethical theory. Whether this word ought to be "claim", "sentence" or "judgment" is not a question I will discuss, and the reader should feel free to substitute in their preferred term if they wish.

²⁸ Note that Kovesi (1967) challenges this idea. Kovesi's views will be discussed in Chapter 3, Section 2, below.

²⁹ Note, however, the underlying assumption that these supposedly paradigmatic descriptive and evaluative claims are representative of their kinds; I will challenge this assumption in Chapter 3, Section 2, below.

³⁰ For an explanation of how a "dichotomy" can turn from a linguistic one into a metaphysical one, see e.g. Putnam (2002: 11-12, 14). The fact that Johnston tries to remain neutral regarding metaphysical questions will be important later when I point out limitations with Lovibond's (1983) criticism of the Fact/Value Dichotomy (see Chapter 2, Section 3, below).

³¹ I discuss Putnam's views in more detail in Chapter 2, Section 2.

term “Fact/Value Dichotomy” as an *umbrella* term covering distinct, but similar views about the relationship between descriptive and evaluative language. In Chapter 2, Section 1, I will explain that Johnston advocates a specific, and all-things-considered less plausible, form of the Dichotomy. What all versions of the Fact/Value Dichotomy have in common, however, *is the idea that there is some sort of fundamental difference in kind between factual descriptive language and evaluative moral language*. It is already clear, then, that Johnston’s view falls under this umbrella. In the rest of this Section, I will give some more details about Johnston’s specific view and his arguments in defence of it.

It is also worth examining some of the implications Johnston draws out from his view that descriptive and evaluative claims are entirely different from one another. He claims that it is a “grammatical point” that “there will always be some input into a moral argument which cannot be verified by science” (1989: 154). He explains that this is a result of the “logical gap between the discoveries of science and the judgements and actions one may derive from them” (1989: 154).

To clarify, it is not that Johnston objects to the idea that good arguments for moral conclusions can sometimes contain empirical, factual premises. For instance, nothing about the following argument

- A) It is wrong to act in a way that will contribute to rendering the earth uninhabitable.
- B) Causing pollution contributes to rendering the earth uninhabitable.
- C) It is wrong to cause pollution.

violates Johnston’s view. He would view A as a moral claim and B as an empirical factual claim. Johnston would insist that there is a premise in this argument that “cannot be verified by science”, namely, the moral claim A. If we accept premise A, then of course finding out the empirical premise B helps guide us to the moral conclusion, C. However, Johnston points out, somebody who disagreed with this argument on moral, rather than

factual, grounds would challenge premise A (1989: 206), and Johnston maintains that moral, evaluative premises like A cannot be proved or supported by appeal to factual, descriptive claims:

There is a desire to claim that our [ethical] concepts and our practices reflect reality or are justified by it; incoherently, we want to claim that what we do is dictated by what is the case, as if our actions themselves could be based on truth (1989: 85).

Why does Johnston think that attempting to justify moral claims by appealing to “reality” – e.g. empirical facts – leads to an incoherence? His answer is that such attempts belie a failure to grasp the “grammatical” distinction between descriptive and evaluative language:

The notion of supporting a moral judgement with evidence of its truth makes no sense because, unlike an empirical judgement, a moral judgement does not assert that a particular state of affairs holds; rather, it asserts that a certain act should be done (or should not be done), hence the notion of evidence for the truth of the judgement gets no grip (Johnston 1989: 201, cf. 1989: 142).

Johnston, then, subscribes to the view I will call the “Humean Claim”:³² that descriptive claims can never be, nor imply, evaluative claims.³³ Johnston holds that descriptive claims can never *be* evaluative claims because the “grammar” of description is of a “fundamentally different kind” to the “grammar” of evaluation, and he holds that descriptive claims can never *imply* evaluative claims either. This means that, as we just discussed, Johnston denies the possibility of validly deducing moral claims from factual, descriptive premises. Moreover, he dismisses the idea that evaluative conclusions might be validly deducible from *conceptual* descriptive premises, too: “no substantive [i.e. moral] claim can be demonstrated (or supported) by purely conceptual argument” (1989: 201). For Johnston, conceptual claims simply delineate the boundary between sense and nonsense;

³² This claim will be explained in more detail in Chapter 2, Section 1.

³³ Like Johnston, I choose to use the word “claim” here, but “judgment” is, in my view, an acceptable substitute. If the reader wishes, he or she can interpret my word “claim” to mean the same as what they call “judgment”.

they say nothing at all about which of the many possible coherent moral views we should take up and endorse as our own (1989: 139).

So far, I have simply been trying to present Johnston's claims as accurately as possible. Why, though, does he believe the Humean Claim?

Part iii: Why Does Johnston Embrace the Fact/Value Dichotomy?

Although I cannot give a fully comprehensive account of Johnston's arguments here, it is important that one of Johnston's main reasons for believing in the Fact/Value Dichotomy is that he thinks it is supported by the right accounts of the semantics of both descriptive and evaluative language. Since, on the accounts he defends, the "grammars" of descriptive claims and evaluative claims *are* utterly different from one another, it is natural enough that Johnston ends up endorsing the Fact/Value Dichotomy. So, to make Johnston's reasons for endorsing that Dichotomy clear, let us consider what he says about the grammar of descriptive language and evaluative language specifically.

Regarding empirical descriptive claims, Johnston writes that "the criteria for their application are fixed; empirical judgements therefore occur in the context of agreed procedures for verification" (1989: 95-6). He claims that these procedures for verification "necessarily" involve a kind of "comparison [between the descriptive claim and] reality" (1989: 142).³⁴ To illustrate this view with an example, consider the descriptive claim "all apples are blue". The "agreed procedures for verification" here involve looking at apples to determine whether they in fact fit with, and verify, the claim or not. Since, of course, they do not, we can conclude on the basis of our "comparison" between reality and the claim – and what it says about apples – that the claim is not verified, and should be rejected.

³⁴ Johnston's view here seems to fall into the category of "pictorial semantics", which Putnam (2002: 15) identifies, and thinks is responsible for the Fact/Value Dichotomy. I will discuss this point further in Chapter 2, Section 2.

In conceiving of descriptive language this way, Johnston falls in line with the tradition of verificationist empiricism which goes back to the likes of Ayer (1936, see e.g. 105-6) and perhaps even further back to Hume (see Putnam 2002: 14-15).³⁵ In line with Lovibond (1983) and Putnam (2002), I will later (in Chapter 3) challenge Johnston's understanding of the meaning of descriptive language, using Julius Kovesi's (1967) work as a springboard with which to develop a novel challenge to the Fact/Value Dichotomy.³⁶

For now, it is important to see that Johnston operates with a conception of the meaning of descriptive statements on which there are "fixed criteria" for their application and verification, which involve a comparison between the claims and reality. Part of Johnston's reason for concluding that the meaning of evaluative claims must be entirely different from this is simply that it is very hard to see how the model that he thinks governs descriptive claims could possibly be applied in the case of evaluative language (see e.g. 1989: 99). Consider, once again, Johnston's contrast between statements such as "this is good" and "this is yellow". When it comes to the second statement, it is straightforward to see how we can look to see whether the statement is correct by simply observing the object in question. Checking whether something is yellow seems to be a matter of simply observing its physical properties. When it comes to goodness, however, it is entirely unclear how an observation of an object's physical properties relates to its goodness, or whether any of them constitutes its goodness.

To illustrate with a specific example, contrast the two statements "there is a red book on this desk" and "eating meat is wrong". When it comes to the first, descriptive claim, we can quite easily imagine a picture of how reality could match with, and verify, the statement. But when it comes to the second claim, how could reality possibly match with or verify the statement in the same way (for discussion see e.g. Putnam 2002: 15)? Since it is hard to make sense of how a picture of reality could verify that "eating meat is wrong", Johnston prefers to conclude that a picture of reality *cannot* verify it, or any other such

³⁵ See Ch. 2, Section 2, below, for more on Putnam's view.

³⁶ See Ch. 3, below.

evaluative moral claim.

This apparent difference between the rules governing our use of descriptive factual claims and evaluative moral claims leads Johnston to think that there is a fundamental semantic difference here (1989: 97), and, in this respect, Johnston is in line with non-cognitivists, who deny that moral claims have the same logical form as descriptive, factual claims, for similar reasons to Johnston (see e.g. Ayer 1936, Stevenson 1945 and more recent non-cognitivists such as Gibbard 1990). We also find the same reasoning in Hume (1985 [1739]: 520).³⁷ For ease of reference, I am going to call to this argument the *Humean Argument*.³⁸

To clarify the “Humean Argument”, I will now summarise Johnston’s reasoning. Johnston has what he takes to be a highly plausible account of the meaning of descriptive claims, which explains the meaning of such claims in terms of a “comparison [between the descriptive claim and] reality” (1989: 142). However, Johnston then finds that he cannot fit evaluative claims, including moral claims, into this mould. If we try to compare a claim like “eating meat is wrong” with reality to check if it is true, it is simply unclear which part or parts of reality we can look to, and such claims hence seem fundamentally different to claims like “there is a red book on the desk”. Hence, whatever the meaning or conceptual nature of moral claims might be, it seems that it cannot be the same as that of empirical descriptive claims. To accept this is to arrive at Johnston’s thesis that there is a fundamental difference in conceptual kind here.

³⁷ Hume writes: “Take any action allow’d to be vicious: willful murder, for instance. Examine it in all lights, and see if you can find that matter of fact, or real existence, which you call *vice*. In which-ever way you take it, you find only certain passions, motives, volitions and thoughts. There is no other matter of fact in the case. The vice entirely escapes you, as long as you consider the object. You can never find it, till you turn your reflection into your own breast, and find a sentiment of disapprobation, which arises in you, towards this action”. Thus Hume believes that, whatever makes it true or right to say that “murder is vicious”, this is not the same sort of thing as what makes it true that “some apples are green”. This leads Hume to consider factual descriptive claims to be logically different in kind from evaluative moral claims, and so it is unsurprising that Hume immediately goes on to challenge whether factual descriptive premises can validly entail evaluative moral conclusions (Hume 1985 [1739]: 521), on the grounds that the evaluative word “ought” “expresses some new relation or affirmation” from those found in factual descriptive premises.

³⁸ I will return to, and reject, the Humean Argument in Chapter 3, Section 2, drawing from and developing ideas from Kovesi (1967). Rejecting the Humean Argument is an important part of my critical response to the Fact/Value Dichotomy, but it does not exhaust it.

As Putnam (2002: 15) says when analysing the very similar argument from Hume, this “Humean Argument” seems quite convincing *if we accept the semantic views which it rests on*. If we agree with Johnston’s characterisation of empirical descriptive language, it looks hard to resist the follow-up claim that moral language cannot fit that mould.³⁹ As I argue later, in Chapter 3, Section 3, exposing precisely what is wrong with this “Humean Argument” clears the way for a more plausible form of meta-ethical Naturalism, which avoids the Fact/Value Dichotomy and solves Parfit’s (2011) “Normativity Problem”.

Part iv: Johnston’s Arguments in Relation to Meta-Ethical Theory

I will now set out how Johnston’s Humean Argument supports certain meta-ethical positions, and what our meta-ethics might look like if we reject it, anticipating my discussion in Chapter 3, below.

The Humean Argument, which appeals to this apparent basic difference between claims like “this is good” and “this is yellow” to support the idea that descriptive and evaluative claims are entirely different in kind from one another, supports the basic non-cognitivist tenet that factual descriptive claims are “cognitive”, in the sense that they are truth-apt and can express beliefs, whereas evaluative claims are not truth-apt and do not express beliefs. According to non-cognitivists, evaluative claims do not state truth-apt beliefs, but rather *express* certain kinds of attitude.⁴⁰ This is one example of an Emotivist view of the kind MacIntyre identified (see Chapter 1, Section 1, above). However, it is not the *only* possible form of “Emotivism”: for reasons I will now explain, Johnston’s view is not non-cognitivist, but it is still Emotivist, in that it denies the possibility of rational resolution of moral debates.

³⁹ However, I will question this further in Chapter 3, Section 2.

⁴⁰ What it *is* to express an attitude, and which attitudes get expressed when we make moral judgments, are further questions, on which different non-cognitivists have disagreed. For an overview of different non-cognitivist expressivist proposals see Schroeder (2008).

Johnston's positive account of the grammar of evaluative claims resembles non-cognitivism, since he claims that evaluative language games "extend and develop certain natural reactions (viz. expressions of approval, condemnation, etc.)" (1989: 96), but he is not a non-cognitivist, since his positive account is more general than that of non-cognitivists. He claims that the key distinguishing mark of evaluative language is its intimate relation to action: "the fundamental contrast here is between investigating reality and deciding how to act" (1989: 142). More clearly:

A moral judgement... unlike an empirical judgement... does not assert that a particular state of affairs holds; rather, it asserts that a certain act should be done (or should not be done) (1989: 201).

Johnston then goes on:

With respect to moral judgements, there is no such thing as correspondence with the facts or with reality. Rather, in this context the claim to truth or objectivity [of some moral judgements] expresses the claim that one set of judgements about how people should act is uniquely correct... the moralist... asserts that one set of standards is correct and should be followed by everyone. ... His claims present us with the need to make a substantive decision and hence one which cannot be made on the basis of empirical evidence or conceptual analysis (1989: 202-203).

So while Johnston does not embrace non-cognitivism, his position is nevertheless Emotivist in MacIntyre's sense (see Ch.1, Section 1, above). For Johnston, there is no question of a moral claim being rationally justified or unjustified. There is just a plurality of evaluative views available, and philosophical reflection does not, and indeed cannot, show us that it is better to choose one way or another. While the "moralist" believes that his views *are* more justified than others, Johnston just takes this to be just a further substantive claim which lacks an "independent foundation" (1989: 203) and can hence be freely rejected at no philosophical cost. Thus, Johnston stands squarely on the side of the Emotivist, who considers the realist project of trying to rationally vindicate certain moral

judgments to be futile.

While the Humean Argument seems to lead to Emotivist views, either in the guise of non-cognitivism or in the guise of Johnston's view, it is also important to note that, as Kovesi (1967: 19) observes, at least one kind of non-Emotivist view also accepts the Humean Argument, namely Moorean Non-Naturalist realism (see e.g. Moore 1903, cf. Parfit 2011). The Moorean agrees that claims like "eating meat is wrong" cannot be made true by their correspondence with states of affairs that can be perceived with the ordinary five senses, in the way that claims like "some apples are green" can be. In that sense, Mooreans allow there is a fundamental difference between moral claims and ordinary, empirical descriptive claims like "some apples are green". However, unlike the Emotivists, Mooreans do not take this to be evidence against the existence of moral truths, because they believe that moral truths are non-natural truths, and that moral properties are non-natural properties. This means that they are properties which we cannot observe through (roughly speaking) "normal" sense-perception, and which do not exist in the natural world, but which are nonetheless real (similarly to how theists typically believe that God exists, but not as part of the natural world).⁴¹

Unfortunately, Moorean non-naturalist realism incurs serious ontological and epistemological problems. As Johnston says, Moore's view raises the question of *how* we come to know about these non-natural properties and truths, if not by ordinary sense-perception. Although Moore invokes the idea of intuition to try to answer this question, "he is unable to specify a vehicle for this faculty, nor does he mention any means by which one might attain agreement as to its deliverances" (Johnston 1989: 97). While this may not, in itself, be a decisive blow to Moorean Non-Naturalism, it does make clear that Non-Naturalism faces some serious challenges: it has to justify its non-naturalist ontology, which includes non-natural moral properties. Moreover, to be epistemologically adequate, it must specify some plausible means by which we can come to have knowledge of these non-natural moral properties. So even if Moorean Non-Naturalism is ultimately defensible, defending it may be a tall order. Opponents of Emotivism may, then, understandably feel

⁴¹ See Johnston (1989: 96-97) for this interpretation of Moore (1903).

that they would not want to pin their hopes of defeating Emotivism on Moorean non-naturalism, and it is partly for this reason that I pursue a non-Moorean response to Emotivism in this thesis (see Ch. 3, below). The ontological and epistemological burdens of Moorean non-naturalist realism, I will argue, are ones we simply do not need to take on (cf. Kovesi 1967, Meynell 1971 and Chapter 3, below).

For the moment, we should note that while Mooreans accept what I called the Humean Argument, they do *not* accept the Humean Claim, because Mooreans think some claims, such as “eating meat is wrong”, are *both* descriptive and evaluative. Claims like this, for the Moorean, are descriptive because they describe non-natural facts, but also evaluative, because they ascribe evaluative properties (e.g. “wrong”), and evaluate particular acts, or kinds of act (e.g. “eating meat”). The non-cognitivist and Johnston, by contrast, deny that such claims *describe* anything (either natural or non-natural facts).

While this means that Moorean Non-Naturalists deny the Humean Claim, I will argue that they still fall under the umbrella of the Fact/Value Dichotomy. It is just that they advocate this Dichotomy in a different form, as we will see in Chapter 2, Section 1 and Chapter 3, Section 3, below.

The reason I considered the Moorean non-naturalist view here is that we have shown that even opponents of Emotivism, like some realists, nevertheless hang on to certain claims and arguments that are part of the Emotivist picture. Since Mooreans accept the Humean Argument, they feel it necessary to proffer an ambitious non-naturalist ontology to hold on to their realism. As I argue later in Chapter 3, Section 2, however, we should reject the Humean Argument. This yields better options for how we can answer Johnston, and the Emotivist, by ensuring we do not concede too much to them.

Another reason this is important is that it suggests the difficult issue may not be in rejecting Johnston’s view, which, I shall suggest, can be done easily enough (see the remaining Sections of this Chapter and Chapter 2, Section 1), so much as *how* and *why* we

disagree with Johnston. The Moorean disagrees with Johnston in a specific way, denying the Humean Claim and advancing a different ontology, but without challenging some of the core arguments supporting Johnston's position, such as the Humean Argument. In a similar way, conventional forms of meta-ethical Naturalism, so I will argue, will reject Johnston's Humean Claim, but still go along with the Fact/Value Dichotomy (see Ch. 3, Section 3, below). So I focus on Johnston's view not so much because I find Johnston's view plausible, but because, in this thesis, I want to identify precisely *what* is wrong with it.

My subsequent argument in Chapter 3 will be that, although the majority of contemporary meta-ethicists would probably disagree with Johnston, they do not do so in the correct way. This is because they embrace an alternative form of the Fact/Value Dichotomy to Johnston. We should, instead, reject⁴² the Fact/Value Dichotomy (much like MacIntyre 1981). Whereas most contemporary meta-ethicists would say that Johnston's view about the relationship between descriptive and evaluative language is *false*, I instead argue that it is *incoherent*. We cannot make sense, I shall argue, of the idea of "pure" descriptive language which is entirely conceptually different in kind from evaluative language, which Johnston appeals to (see Ch. 3, Sections 1 and 2, below). Therefore, engaging with Johnston's view (even if Johnston's view is implausible) delivers important results.

To recap, Johnston subscribes to the Fact/Value Dichotomy, specifically the Humean Claim, which is that a descriptive claim can never be, nor imply, an evaluative claim. He does so because, on the accounts of descriptive and evaluative language he accepts, the grammar of these two sorts of language are of "fundamentally different kinds" (1989: 96). Johnston understands the meaning of descriptive language in terms of a comparison between the claim and reality, but finds he cannot apply the same model to evaluative language, which gives him the impression that there is a sharp distinction between the two kinds of language.

⁴² I will clarify what I think is involved in rejecting the Dichotomy in Chapter 2 and Chapter 3, Section 1.

Part v: Johnston's Insistence on Ethical Neutrality

Having highlighted the place of the Fact/Value Dichotomy in Johnston's position, I now want to explore Johnston's insistence that his account is ethically neutral. This will help us understand why Johnston's position might seem attractive and plausible. Later, though, I argue that Johnston's account cannot deliver on this attractive quality of ethical neutrality (see Chapter 1, Section 3, below).

To explore this, I need to introduce some terminology Johnston uses. Johnston calls evaluative ethical claims which concern e.g. what is morally good, morally bad, morally wrong, etc. "substantive" claims. According to Johnston, these are to be distinguished from "grammatical claims" *about* ethics, which do not speak of what is morally good or bad, but are rather about ethical language. So, for instance, to say "the hallmark of an ethical statement is its intimate relation to action" is not to make a "substantive" ethical claim like "eating meat is wrong"; it is rather to make a descriptive statement *about* ethical language.

In line with this purported distinction between "substantive" and "grammatical" claims, Johnston distinguishes between "first-order" enquiry in moral philosophy, which involves advancing and arguing for substantive ethical claims ("normative ethics" is the standard term, see e.g. Schroeder 2017) and "second-order" enquiry in moral philosophy, which aims to clarify moral language. Although Johnston does not himself use this term, many philosophers would call second-order enquiry of this sort "meta-ethics". (While meta-ethics is typically thought to encompass *more* than just clarification of moral language (since it is also typically taken to include provision of an epistemology and ontology of morals), the clarification of moral language is usually considered to at least be in the remit of meta-ethicists). So Johnston's "second-order" enquiry is recognisably a kind of meta-ethics.

As mentioned in part i) of this Section, above, Johnston considers his investigation to be purely descriptive, and so it is no surprise that he sees his own investigation as "second-

order”. What is more surprising – and, as we shall see,⁴³ controversial – is that he insists that he can and must remain entirely neutral on all first-order matters in giving his second-order account. For Johnston, “the grammatical investigation of ethics... is non-substantive in nature” (1989: 209-10). There are two main reasons why Johnston comes to this view.

First, since Johnston believes there is a sharp distinction between descriptive and evaluative language, it coheres with his general account to say that there is also a sharp distinction between descriptive and evaluative moral philosophy.

Second, as we saw in part ii) of this Section, Johnston maintains that “no substantive claim can be demonstrated (or supported) by purely conceptual argument” (1989: 201). For Johnston, his approach simply involves elucidating the meanings of concepts. From his point of view, his descriptions of moral concepts simply tell us what, logically, *can* be said about morals (i.e. what can be said without lapsing into nonsense) and are silent about what we *should* say and do. Therefore, he thinks that his method, if carried out faithfully, ends up being silent about evaluative matters (see e.g. 1989: 210). Since he thinks conceptual claims cannot support or justify substantive claims, he concludes that mixing together grammatical and substantive enquiry is, at best, a recipe for confusion, and, at worst, a deliberate attempt to motivate one’s substantive views by illicit means (see e.g. 1989: 160).

One reason that Johnston’s stance on neutrality is important for us to consider is that Johnston presents the supposed neutrality of his position as being one of the main advantages of his view over rival meta-ethical approaches. He argues that other main meta-ethical positions involve committing to certain normative ethical claims.⁴⁴ One example, which I mentioned in part iv) of this Section, is his treatment of realism: for Johnston, the realist’s belief that some moral claims are more justified than others is simply a *further*

⁴³ In Chapter 1, Section 2.

⁴⁴ For Johnston, this means that these rival views are a confused mixture of descriptive and evaluative material. He takes one of the advantages of his own view to be that it sticks rigorously to description, and not evaluation (see 1989: Conclusion).

“substantive” claim which lacks an “independent foundation” (1989: 203). To my mind, Johnston’s claim here is at least somewhat plausible: the realist’s view that some moral claims are more justified than others seems to be contestable on evaluative grounds. To deny the realist’s claim that “some moral claims are more justified than others” would be to maintain that no moral claim is more or less justified than any other, and that seems like an evaluative, rather than a purely descriptive, thesis. Hence, according to Johnston, the realist’s position, unlike his own, depends on a controversial normative ethical claim. (As it happens, Johnston also thinks that many non-realist meta-ethical theories depend on substantive ethical claims (see e.g. 1989: 160, where Johnston says this applies to Mackie’s (1977) error theory), but there is no need to go into this in detail here, since it is primarily the dialectic of “Johnston vs the realist” that I am focusing on in this thesis). I am happy to grant Johnston that realist meta-ethical positions depend on controversial normative ethical claims (see the conclusion of Chapter 1, Section 3, below, for more on why I am happy to allow this).

Johnston believes that the (supposed) ethical neutrality of his position gives it an advantage over rival approaches. As Mulhall (2002: 302) comments, if Johnston can make good on his claim to ethical neutrality, it would seem to give his account a special kind of “philosophical authority”, since (unlike its realist rivals, Johnston will say) it does not depend on any contentious normative ethical claims, and can hence be accepted by any philosopher irrespective of his or her normative ethical views.

In addition to his direct arguments in support of the Fact/Value Dichotomy, then, Johnston also believes his position has another significant merit: its ethical neutrality. Exactly how large an advantage this would be over non-neutral views depends, of course, on how controversial the normative ethical claims are which non-neutral views need to accept, but it nevertheless seems attractive for a meta-ethical theory not to need to get embroiled in normative ethical debate. This is especially so considering the high level of controversy around whether it is rationally possible to resolve normative ethical debates (on which see Chapter 1, Section 1, above). While realists may consider it unproblematic for their position to depend on certain normative ethical claims, Johnston and other Emotivists will disagree, because they do not think that such claims can be rationally justified. So,

Johnston could say (reasonably enough), it is dialectically infelicitous for realists to depend on normative ethical claims to defend their realism - at least if they are hoping to win round any Emotivists.

For that reason, I think it is important to critically assess whether Johnston's account (and, indeed, other would-be ethically neutral meta-ethical theories) can make good on this claim to ethical neutrality. In the next Section, I shall argue that they cannot. Later in the thesis, in Chapter 3, Section 2, I will also tackle Johnston's direct argument for the Fact/Value Dichotomy.

Conclusion to Chapter 1, Section 2

I will now recap what I have done in this Section. First, in part i), I outlined Johnston's methodology, which is based on his distinction between the description of our moral language and arguing for "substantive" normative ethical claims about how we morally ought to act. We saw that Johnston wishes to stick strictly to the first of these activities and avoid the second. In parts ii) and iii), I summarised Johnston's argument for the view that descriptive and evaluative language are, in his words, of "fundamentally different kinds" (1989: 95-96), and then in part iv) I discussed his view in relation to other meta-ethical theories, looking especially at how other meta-ethical theories such as Moorean realism agree with Johnston on certain points even while differing from him on others. Finally, I discussed Johnston's claim that his account is ethically neutral, and why he thinks this is an advantage of his account.

Although I do not find Johnston's arguments convincing, for reasons I will explain in Chapters 2 and 3, it will be helpful to understand, and engage with, the (apparent) advantages of Johnston's position, even if we do not find Johnston's view plausible *prima facie*.

First, Johnston's "Humean Argument", which uses the apparent contrast between the meaning of words like "good" and words like "yellow" to motivate the idea that there is a fundamental logical difference between descriptive and evaluative language, is widely thought to be plausible. On the basis of this argument, Hume famously asserts that a claim about what we "ought" to do (which is evaluative) "expresses some new relation or affirmation" to a claim about what "is" the case, and that it "seems altogether inconceivable, how this new relation [ought] can be a deduction from the others ["is" claims], which are entirely different from it" (Hume 1985 [1739]: 521). As we will see in Chapters 2 and 3, many meta-ethicists have been much-influenced by this Humean idea, even those who do not advocate the stronger form of the Fact/Value Dichotomy Johnston embraces. (For further discussion and examples of this see MacIntyre 1959; my argument in Chapter 3, Section 3 will be that Parfit (2011) accepts too much of this Humean picture of descriptive language and evaluative language being "entirely different" in kind, in his meta-ethics).

So, the first apparent advantage of Johnston's account is that, on the face of it, Johnston simply *gets it right* about the nature of descriptive and evaluative language: we might be inclined to agree that there *is* a "fundamental difference in kind" here. Like Hume, perhaps, we might just think that appearances suggest that this difference in kind is present, and rest content to leave things there. Since this is the central claim in Johnston's account, we might judge that Johnston's overall package of views carries some appeal. At the very least, I suggest, opponents of Emotivism need to find the best way to answer Johnston's arguments, which is a task I take up in Chapter 3, below.

The second apparent advantage of Johnston's account is that he (purportedly) remains neutral on all substantive ethical matters. If Johnston can make good on his claim to ethical neutrality, that would mean that any philosopher can, in principle, accept his account, *irrespective* of their first-order ethical views. This would insulate Johnston's account from criticisms made for first-order reasons and for that reason, as Mulhall writes, Johnston's

account seems to possess a kind of “philosophical authority” (2002: 302).⁴⁵ I shall argue in the next Section, however, that Johnston cannot make good on his claim to ethical neutrality.

The fruitful philosophical results to be yielded here are not so much to do with proving *that* Johnston’s core claims are wrong, but rather to understand *why* Johnston is wrong. As we will see, my way of explaining *why* and *how* Johnston is wrong differs from the reasons that most meta-ethicists, like Parfit for example, would give (see Chapter 3, below).

⁴⁵ Arguably, other kinds of meta-ethical positions, such as non-cognitivism, share the same advantage (see e.g. Schroder 2016). But this would remain an advantage of Johnston’s view, even if it is not one that Johnston’s view enjoys exclusively.

Chapter 1, Section 3: Can There Be An Ethically Neutral Meta-Ethics?

As I mentioned at the end of Section 2, Johnston (1989) claims that his meta-ethical account is neutral regarding normative ethical matters, and argues that this gives his account an important advantage over rival accounts, including realism. In this Section, I argue that Johnston cannot make good on this claim to ethical neutrality. In fact, I shall argue that there cannot be an ethically neutral meta-ethics. If this is right, then clearly Johnston and other Emotivists cannot rightly claim ethical neutrality to be a comparative advantage of their accounts over realist or other meta-ethical accounts.

If the reader is less than fully convinced by the argument of this Section, the implications for the overall argument of this thesis are not too grave. Firstly, if the reader is not fully convinced of the impossibility of an ethically neutral meta-ethics, the arguments of this Section can at least be taken to *motivate scepticism* about whether a meta-ethical account like Johnston's can really be entirely neutral. Secondly, my main critical response to Johnston, which I develop in Chapters 2 and 3, targets the Fact/Value Dichotomy. The argument of this Section does not directly bear on my argument against the Fact/Value Dichotomy (although there is an indirect connection, which will be explained in the next Section of this Chapter).

Nevertheless, the argument of this Section is important, because it shows that Johnston's Emotivist view does not, in fact, hold one of the main advantages over realism that he thinks it does. At the end of the last Section, I explained why, *if* Johnston was right that only his account (and not the realist's) is ethically neutral, this would put the realist in a disadvantageous dialectical position. The reason is that Emotivists challenge whether normative ethical claims can be rationally justified, so if a realist appeals to normative ethical claims in their arguments against Emotivism, Emotivists are unlikely to find their arguments convincing.

It is instructive again to consider the contrast between Johnston and MacIntyre on the issue of ethical neutrality of meta-ethics. As we saw in Chapter 1, Section 2, part iii), Johnston insists on the strict ethical neutrality of his meta-ethics, which fits with his idea that description and evaluation are logically separate. MacIntyre, by contrast, claims that “descriptive” meta-ethics *cannot* be entirely ethically neutral, which fits with his idea that description and evaluation necessarily go together (see MacIntyre 1959: 3 and MacIntyre 1981: ix). In this Section, I argue that MacIntyre is right about this: there cannot be an ethically neutral meta-ethics. Then, in the next and final Section of this Chapter, I will pinpoint which normative ethical claims are inconsistent with Johnston’s meta-ethical position, helping us identify a route of attack against Johnston’s Fact/Value Dichotomy.

I will explain how my rejection of the possibility of an ethically neutral meta-ethics runs contrary to the prevailing ‘Neutrality Thesis’ (endorsed, for instance, by Schroeder (2017)), according to which some, but not all, meta-ethical positions are ethically neutral (see part i, below). I will then outline Mulhall’s argument against the Neutrality Thesis in parts ii) and iii). Finally, in parts iv) and v), I identify why Mulhall’s argument does not work, and offer a refined version of his argument which avoids the problems faced by Mulhall’s original argument while still showing the Neutrality Thesis to be false.

Part i: The Case for the Neutrality Thesis

Moral philosophers often say there are two kinds of claims we can make about morality. First, there are *normative ethical* claims about what is morally right and wrong, permissible and impermissible, etc. For example, one might claim that telling the truth is morally right. Secondly, though, we can make claims about what it *means* to say things like ‘telling the truth is morally good’.

Corresponding to this intuitive distinction are two apparently distinct branches of moral philosophy: *normative ethics*, which deals with what is morally right and wrong,

permissible and impermissible, etc., and *meta-ethics*, which deals (among other things) with the question of *what it is* to judge an act or practice to be morally right or wrong (Schroeder 2017). The classic debate between meta-ethical realists (e.g. Moore 1903, Parfit 2011) and non-cognitivist expressivists (e.g. Blackburn 1984, Gibbard 1990), for example, focuses on *what is involved* with asserting a normative ethical claim, like ‘telling the truth is morally good’, and is hence considered a meta-ethical debate (cf. Ch. 1 Section 1, above). Does the person uttering such a sentence express a truth-apt claim much like the person who utters ‘the grass is green’, as the realist claims? Or are they expressing a certain emotive attitude towards truth-telling, as expressivists argue?

In highlighting the connection between meta-ethical enquiry and language, I do not assume that *all* that is involved with meta-ethics is the analysis of moral language. Meta-ethicists also often analyse epistemological and ontological issues relating to ethics, for example. I do assume, however, that analysis of moral language is a *necessary part* of meta-ethics: any enquiry which warrants the title ‘meta-ethics’, I believe, must include at least some such analysis.⁴⁶

In this Section, I discuss the relationship between normative ethics and meta-ethics. According to a widespread and seemingly plausible view I call the *Neutrality Thesis*, which is explicitly defended by Schroeder (2017), it is possible to make at least some meta-ethical claims without thereby committing oneself to any normative ethical claims. Thus, on this view, it is possible for one’s meta-ethical position to be *ethically neutral* in the sense that it does not carry any normative ethical commitments (in other words, it does not contradict any normative ethical claims).

⁴⁶ I will not be defending this view of meta-ethics here, since it is taken for granted in the main literature I will discuss, but if the reader disagrees with me on this, they should feel free to consider the scope of my argument differently. I believe that my argument shows that there cannot be an ethically neutral meta-ethics, and that meta-ethics necessarily involves analysis of moral language. If the reader believes that meta-ethics does not necessarily involve analysis of moral language, they should understand my argument to support the more limited conclusion that any meta-ethics *which includes the analysis of moral language* cannot be ethically neutral.

Although I ultimately reject the Neutrality Thesis, I accept that it seems plausible *prima facie*. I now want to set out why philosophers like Schroeder find the Neutrality Thesis so plausible, which will help set up my argument against it.

What would be wrong with the view that *all* meta-ethical claims are neutral regarding normative ethics? As we saw in Chapter 1, Section 2, some philosophers, such as Johnston, seem to think the answer is “nothing”. In a similar vein, Ayer writes that meta-ethicists ‘attempt to show what people are doing when they make moral judgments... *not* [give] a set of suggestions as to what moral judgements they are to make’ (Ayer 1972: 245-6, emphasis mine). While Johnston thinks that some meta-ethical theorists, like realists, advance both normative ethical claims and meta-ethical claims (see Section 2, part v), above), Johnston sees this as a transgression beyond their proper place as meta-ethicists. That is, for Johnston, if these theorists confined themselves to only making meta-ethical claims (as, in his view, they should do), they would have had no need to make or commit to any normative ethical claims, since he thinks all meta-ethical claims are ethically neutral.

However, the view that all meta-ethical claims are neutral regarding normative ethics is now widely perceived to be false (see e.g. Schroeder 2017, Christensen 2019). Schroeder gives apparently decisive counter-examples to the view that meta-ethical claims are all neutral regarding normative ethics. Schroeder’s first example is an error-theoretical position combining the claims

- i) if any actions are morally required, there must be categorical reasons to do those actions, and
- ii) there are no categorical reasons (see Schroeder 2017: 677).

An error theorist embracing i) and ii) must conclude that no actions are morally required, but this is inconsistent with all normative ethical views which say that at least some actions are morally required (Schroeder 2017: 677). Since i) and ii) appear to be meta-ethical

claims, it seems that Schroeder has found an example of a pair of meta-ethical claims which entail a normative ethical conclusion.

Schroeder's second example is of certain naturalist meta-ethical views. Consider, for example, the reductive naturalist view on which the property of moral rightness just is the natural property of maximising happiness (Schroeder 2017: 681). This thesis, which appears to be a meta-ethical one, straightforwardly entails that all acts which maximise happiness are morally right. That, however, is inconsistent with various normative ethical claims and theories according to which moral rightness and maximising happiness sometimes come apart (Schroeder 2017: 682-3).

Schroeder's view, of course, is a middle-way between the extremes of Johnston's view, on which all meta-ethical positions are ethically neutral, and MacIntyre's view, on which no meta-ethical positions are ethically neutral. I believe Schroeder's counterexamples are convincing, and prove that, *pace* Johnston, at least some meta-ethical claims include normative ethical commitments (and are hence non-neutral). To resist Schroeder's counterexamples, Johnston would have to dispute whether the error-theoretical and naturalist claims which feature in Schroeder's counterexamples are genuinely *meta-ethical* claims. Johnston construes meta-ethics in a particularly narrow way, arguing that meta-ethics is restricted to mere 'grammatical' claims about language (see Mulhall's comment on Johnston at 2002: 299), and this gives him a basis for believing that views like error theory and naturalism are hybrids of meta-ethical and normative ethical views (see e.g. Johnston 1989: 159-60). He would hence argue that genuine, 'pure' meta-ethical claims are entirely neutral regarding normative ethics (see Chapter 1, Section 2, above). However, I believe Johnston's view is highly revisionary, and over-restrictive. Error theory and naturalism are paradigmatic meta-ethical views, and any view that says otherwise is, I suggest, implausible.

If the Neutrality Thesis holds, then there is no reason in principle to think that Johnston's account cannot be ethically neutral. Moreover, even if we find Johnston's particular Emotivist account not to be neutral, realists may have a lingering worry that some

alternative Emotivist account could be provided which *does* enjoy neutrality, even if Johnston's does not.⁴⁷ (Schroeder, for instance, identifies "traditional non-cognitivism" as an ethically neutral meta-ethical view (2017: 676) and which by our classification is also an Emotivist view, so this is not an idle concern). For the reasons we have discussed, this would put realists at a dialectical disadvantage *vis-à-vis* Emotivism. It would hence be good news for realists if the Neutrality Thesis turned out to be false, as MacIntyre and Mulhall believe it is.

Since Schroeder's counter-examples are persuasive, it is safe to say that at least *some* meta-ethical claims are inconsistent with neutrality. Why, then, not go even further and say that *all* meta-ethical claims carry normative ethical commitments? Most philosophers do not want to go so far, instead settling for Schroeder's "Neutrality Thesis", according to which it is possible to have an ethically neutral meta-ethics, even though there are some non-neutral meta-ethical positions.

The main consideration supporting the Neutrality Thesis is that it is hard to see why meta-ethical premises would *necessarily* entail normative ethical conclusions (Schroeder 2017: 675-676). For example, if I begin to believe the non-cognitivist meta-ethical view that moral judgments do not express beliefs, why would this commit me to any normative ethical views about what is right and wrong? It seems that my first-order moral views would all remain unaffected. Only my understanding of what holding those normative ethical views *amounts* to would have changed, not the views themselves. So at least *some* meta-ethical views, like, perhaps, Emotivist non-cognitivist ones, seem free of normative ethical baggage.

Hence, any successful argument against the Neutrality Thesis would need to motivate the idea that meta-ethical claims *necessarily* have normative ethical implications while also explaining why certain meta-ethical views (such as non-cognitivism) *appear* to be neutral.

⁴⁷ Indeed, I will argue, in the next Section, that Johnston's position is not ethically neutral, since it conflicts with at least some versions of Classical Teleological normative ethics.

I will try to further clarify what it takes to make a successful argument against the Neutrality Thesis in the next part of this Section, before moving on to discuss Mulhall's (2002) argument and then advance my own argument against the Neutrality Thesis.

Part ii: A Strategy for Undermining the Neutrality Thesis

If we are to conclude that an ethically neutral meta-ethics is not *possible*, we firstly need to clarify which kind of possibility we are talking about. To clarify, I am focusing on *conceptual* possibility (and impossibility) here. My point is that any claim or position which genuinely warrants the name of *meta-ethical* cannot be ethically neutral. In other words, any claim or position which genuinely fits with our *concept* of 'meta-ethical' cannot be ethically neutral. So rather than involving a substantial metaphysical commitment, for instance, my opposition to the Neutrality Thesis is based on what I take to be an accurate understanding of the concept of 'meta-ethics'.

To support the idea that meta-ethical claims necessarily come part and parcel with normative ethical commitments, the most promising path is to argue that there are certain aspects of making a meta-ethical claim which necessarily involve making some normative ethical claim or judgment. In plain language, we should identify some 'hurdles' which one must cross if one is to count as making a meta-ethical claim at all. If such a 'hurdle' can be identified, we then need to examine whether crossing that hurdle necessitates committing oneself to one or more normative ethical claims. If it does, then the Neutrality Thesis is false, because in that case one must commit to at least one normative ethical claim if one is to so much as *make* a meta-ethical claim at all. This is roughly the strategy followed by Mulhall (see Mulhall 2002: 299-302). I will look at the details of Mulhall's argument in the next part of this Section.

From this, the reader can discern the criterion for ethical neutrality I am working with in this Section: a meta-ethical claim or position is ethically neutral, I say, if and only if it is consistent with all normative ethical claims. This is a fair criterion, because if a meta-ethical claim is inconsistent with at least one normative ethical judgment, it would be absurd to say that it is ethically neutral, since it *rules out* at least one normative ethical judgment. That would mean that somebody who accepted that meta-ethical claim could not remain entirely neutral in normative ethics, because their hands would be tied with respect to at least one normative ethical judgment. My use of this criterion also follows previous work on this topic by Schroeder (2017), Johnston (1989) and Mulhall (2002).

The first step of my argument must then be to identify the ‘hurdle’ which we must cross if we are to count as making a meta-ethical claim at all.

Recall that, as I mentioned above, on my view meta-ethics necessarily involves (among other things, perhaps) analysis of moral language. To be more specific, I agree with Mulhall that, to make a meta-ethical claim, we must make some sort of decision about what does and does not count as moral language (see Mulhall 2002: 299). Take, for instance, the famous non-cognitivist claim that moral judgments do not express beliefs. For this claim to be meaningful and substantive, we would need the non-cognitivist to give us some sort of story about what moral judgments are. Otherwise, how are we to know *which* judgments, exactly, she is saying do not express beliefs? A similar point applies, *mutatis mutandis*, to the realist and error-theoretical claim that moral judgments are truth-apt.

This shows that, to echo Mulhall (2002: 302-303), in the absence of some sort of (implicit or explicit) rationale for distinguishing moral from non-moral judgments, it would be entirely unclear whether the person saying ‘moral judgments do not express beliefs’ would really be analysing *moral* judgments, or whether they were merely under the illusion of doing so (and in reality analysing some quite different, narrower or broader, kind of judgment and/or language). Having such a rationale is not *sufficient* to ensure one is making a genuine meta-ethical claim or claims, since one’s rationale may be ridiculously implausible. But I think Mulhall is right that having such a rationale is *necessary* for

successfully making a meta-ethical claim. Without one, there is no reason to believe that one is focusing on the subject-matter of meta-ethics, namely moral phenomena (including, especially, moral claims (or moral judgments, if the reader prefers)).

I have now identified the ‘hurdle’ which we must cross if we are to make meta-ethical claims: we must have something to say about what distinguishes a moral judgment from a non-moral judgment. There is no reason to assume that what we say here must be so detailed as to allow us to always perfectly discriminate between moral and non-moral judgments. We may sometimes make mistakes in this regard, and our distinction may be rough. But if we have nothing at all to say about this, what sense would there be in saying ‘moral judgments do not express beliefs’? Such a claim is only substantive (in the ordinary sense, not Johnston’s use of “substantive” to mean “normative ethical”) and informative alongside a background idea of what does and does not count as a moral judgment (however rough that idea may be). We cannot work backwards, for example by stipulating that any language or judgment which does not express a belief is a moral judgment. For one thing, this would make our non-cognitivist claim viciously circular. Secondly, it is obvious that many sentences which do not express beliefs are not moral judgments, so this approach would be a non-starter in any case.⁴⁸

The distinction between moral and non-moral judgments, and moral and non-moral language, is controversial. Diamond, for example, is sceptical about the idea of ‘non-moral language’ because, as she shows, sometimes language that would typically be considered “non-moral” can be used to convey moral insights (Diamond 1996, cf. Mulhall 2002). However, I think Harcourt (2019) has adequately answered Diamond on this point: the (undisputed) fact that seemingly non-moral language can be used to convey moral insights is perfectly compatible with the very plausible idea that some language is, nevertheless, distinctively moral. If I describe somebody as just and honourable, there can be little doubt that I have used moral language in my description. Part of the task of meta-ethical theory,

⁴⁸ E.g. if I say ‘I like you!’, I am not expressing a belief, but neither am I making a moral judgment. A similar point would apply for an approach saying that truth-apt claims are all moral claims: there are clear counter-examples to this, and it would make realism viciously circular.

in my view, is to give an accurate account of this kind of language, a task that I intend to take up in Chapter 3, Sections 2-3.

Since, then, I think it is defensible to make some kind of distinction between moral and non-moral language (even if not a very sharp one), and indeed I think meta-ethicists need to have something to say about this distinction if their meta-ethical position is to avoid emptiness, we have identified one possible junction at which meta-ethicists may be forced to give up their neutrality: in giving their account of the very distinction between moral and non-moral language. This is exactly the idea which Mulhall builds his argument against the Neutrality Thesis around, and rightly so, as I will argue later in this Section. I will now lay out Mulhall's argument, so we can evaluate what he gets right and what he gets wrong in his argument against the Neutrality Thesis.

Part iii: Mulhall's Argument Against the Neutrality Thesis and MacIntyre's Moral/Manipulative Distinction

Mulhall suggests that the need to say something about the difference between moral and non-moral judgments forces would-be meta-ethicists into committing themselves to at least some normative ethical claims. For, Mulhall says, the issue of how to distinguish moral from non-moral judgments is itself a normative ethical issue: 'there is no way of characterising the subject-matter of moral philosophy that will not itself give expression to one's own ethical interests and concerns, and that can accordingly be challenged by those who do not share those interests and concerns' (2002: 303). Moreover, 'the distinctions and contrasts without which [the concept of morality] loses its character are not morally neutral' (2002: 303). While (as we shall see in parts iv and v, below) my specific way of arguing against the Neutrality Thesis is different from Mulhall's, I believe Mulhall is right to identify that this is why the Neutrality Thesis fails.

Why, though, would our decision about how to distinguish moral from non-moral judgments commit us to any normative ethical views? The reason I will be focusing on is

that this distinction is ethically controversial because taking a stance on it inevitably involves taking a stance on the distinction between moral and manipulative language, and this distinction is controversial for normative ethical reasons (see Mulhall 2002).

What Mulhall tries to prove, then, is that would-be meta-ethicists must (to make any meta-ethical claims at all) commit to at least one normative ethical claim, because they must take a non-neutral normative ethical stance on the distinction between moral and manipulative language. To explain this issue about the distinction between moral and manipulative language and how it ties in to the distinction between moral and non-moral language, Mulhall refers to the disagreement between MacIntyre (1981) and Cavell (1979), on the one hand, and Stevenson (1945), on the other.

First, though, we should ask: What *is* the purported distinction between moral and manipulative language? For Mulhall, to understand this distinction we need only consider the difference between two ways of trying to persuade somebody to help you. Normally, we might say something like ‘it would be morally right for you to help me out!’, perhaps because we are trying to draw attention to (what we take to be) a genuine *reason* for the person in question to help us. This seems like an example of a use of *moral* language. By contrast, you might also try to persuade somebody to do business with you by saying ‘help me out – or else!’. In this case, rather than trying to identify a moral reason for one’s interlocutor to do business with you, you would be trying to manipulate them into doing business with you. It is natural, then, to call this kind of language manipulative language.

It is this apparently common-sense distinction which the debate between Stevenson and his critics, Cavell and MacIntyre, centres around, and which, Mulhall argues, reveals the impossibility for a neutral meta-ethics. To understand this debate, consider first the highly inclusive way of conceiving of moral language offered by Stevenson (1945). For Stevenson, Mulhall notes, ‘the only legitimate limit on counting any consideration as an ethical one is that it have some possible link with patterns of practical activity’ (2002: 301). If we applied this Stevensonian idea to the question of what differentiates moral from

non-moral language, then, we would conclude that *any* language can be moral, so long as it has possible links with patterns of practical activity.

Perhaps the most natural reaction to the Stevensonian view is that it is *overly* inclusive. Many kinds of language have possible links with patterns of practical activity, not just moral language. So if the Stevensonian view is supposed to capture something *unique* to moral language, it surely fails. Mulhall, however, draws attention to the fact that this Stevensonian view on the difference between moral and non-moral language is objectionable not only because of its over-inclusivity (which is, so to speak, a meta-ethical fault) but also on normative ethical grounds.

Mulhall explains that MacIntyre's reason for objecting to the Stevensonian view on normative ethical grounds is that he believes it 'obliterates' the distinction between moral and manipulative language (and, therefore, the related distinction between moral and manipulative social relations) (MacIntyre 1981: 23-4), and this, MacIntyre suggests, has deeply concerning social consequences for the societies which lose the distinction (1981, Ch. 3). To put MacIntyre's point simply: if we consider manipulative behaviour and utterances to be cause for moral concern, it is vital that we have some way of distinguishing between manipulative utterances and moral utterances, and manipulative social relations and moral social relations. Without these vital tools, we will struggle to differentiate between these things, with the inevitable result that it will become more difficult for us to identify, and critically respond to, instances of manipulative social relations and the language used to sustain them.

To clarify MacIntyre's objection, consider the following, related objection he makes against Stevenson's (1945) emotivism. According to Stevenson, to say 'X is morally good' is roughly equivalent to saying 'I approve of X; do so as well!'. MacIntyre, however, is dissatisfied with this analysis of the meaning of 'X is morally good', on the grounds that it seems to equate a *moral* statement ('X is morally good') with a statement that is merely manipulative ('I approve of X, do so as well!'). When making the latter statement, rather than saying that there is some genuine valid *reason* for approving of

X, the speaker (MacIntyre thinks) simply urges his interlocutor to approve of X because that is what he wants her to do. The statement 'X is good', by contrast, can (for MacIntyre), at least sometimes, be used to draw attention to a genuine valid *reason* to approve of X, one that is not reducible to a given person *wanting* you to approve of X (MacIntyre 1981: 23-4). For MacIntyre, this shows that 'X is morally good' and 'I approve of X, do so as well!' should not be assimilated, as Stevenson thinks they should be, since only the former has a genuine claim to be an instance of moral language.

MacIntyre's idea is that, if we disapprove of manipulative behaviour and social relations on moral grounds (as, surely, many of us do), the overly-inclusive Stevensonian view is troubling, not only because it seems to fail to accurately capture what moral language *is* (a meta-ethical fault), but also that it (if widely accepted) would seem to undermine our ability to recognise, and critically respond to, manipulative behaviour. But undermining that ability, MacIntyre opines, will have socially and morally regrettable consequences. This is why there is a normative ethical basis to MacIntyre's objection to Stevenson: for MacIntyre, accepting Stevenson's view carries a moral cost, in view of the social consequences it is likely to have.

That MacIntyre's concern with the Stevensonian view is normative ethical in kind should not be that surprising, or controversial. As he emphasises (MacIntyre 1967: 34, 1981: Ch. 3), our understanding of our concepts, and the conceptual distinctions we do (and do not) draw, directly affects our actions and practices: 'only those... decisions are available to us which there are concepts available to express' (1967: 34). For MacIntyre, then, it is natural to scrutinise the practical social consequences of the Stevensonian view (and others like it). Moreover, MacIntyre objects to the Stevensonian due to what he takes to be its expected and, in his view, negative social consequences. Such a pattern of argument is a customary way of making a normative ethical objection against some view or practice, and will, for instance, be familiar to all consequentialists. This is why I consider MacIntyre's objection to Stevenson to be at the normative ethical level (as well as meta-ethical).

For MacIntyre, then, there is a *requirement* on any would-be meta-ethical view to respect the distinction between moral and manipulative language. For MacIntyre, as I have explained in the last few paragraphs, any view which fails to respect that distinction will, like the Stevensonian view, be objectionable both on conceptual grounds and for normative ethical reasons. For ease of reference, I am going to call this *MacIntyre's Requirement*. If a meta-ethical view respects the distinction between moral and manipulative language, I will say it satisfies MacIntyre's Requirement. If it fails to do so, I shall say it does not satisfy MacIntyre's Requirement. The Stevensonian view that any language can be moral if it has 'a possible link to patterns of practical activity' is one example of a view that does not satisfy this Requirement, because it includes both manipulative and moral language in the same category.

To take stock: MacIntyre's objection to the Stevensonian way of delineating between moral and non-moral language appears to be based partly on a meta-ethical, conceptual concern, but also partly on a normative ethical concern. Although I have not assessed whether MacIntyre's objection to the Stevensonian view is valid, we can now understand why the Stevensonian view is controversial for normative ethical reasons, even if we do not share MacIntyre's concern about the Stevensonian view.⁴⁹

In the course of outlining MacIntyre's objection to Stevenson, I explained *MacIntyre's Requirement*, which is the requirement for a meta-ethical view to respect the distinction between moral language and manipulative language. "MacIntyre's Requirement" plays an important role in the arguments of both Mulhall (2002) and the present author against the Neutrality Thesis.

Mulhall's idea is that whether, and how, we distinguish between moral language and manipulative language has serious implications for broader social practices, such as the recognition of and responses to manipulative behaviour (such as the issuing of threats). Since, Mulhall thinks, what we say about the difference between moral and non-moral

⁴⁹ As I make clear in part v of this Section, the argument of this Section does not depend in any way on MacIntyre being *right* about all this, unlike Mulhall's argument.

language bears directly on whether and how we distinguish between moral and manipulative language, what we say about the former will inevitably be controversial for normative ethical reasons – as the example of Stevenson’s view illustrates. In combination with the idea that any meta-ethicist must have something to say about the difference between moral and non-moral language, this gets us to Mulhall’s conclusion that there cannot be an ethically neutral meta-ethics.

More formally, we can set out Mulhall’s argument as follows:

Q1) A meta-ethical position must satisfy MacIntyre’s Requirement.

Q2) A meta-ethical position which satisfies MacIntyre’s Requirement cannot be ethically neutral.

C) A meta-ethical position cannot be ethically neutral.

Mulhall indicates his agreement with MacIntyre’s criticism of Stevenson (hence **Q1**) when he writes that Stevenson’s view, by threatening the distinction between moral and manipulative language, “threatens the obliteration of the concept [of moral language] supposedly under analysis” (2002: 303), but also acknowledges that “the distinctions and contrasts without which that concept loses its character are not morally neutral” (2002: 303) (hence **Q2**). “The conclusion to draw”, he continues, is that “there is no way of characterizing the subject-matter of moral philosophy that will not itself give expression to one’s own ethical interests and concerns, and that can accordingly be challenged by those who do not share those interests and concerns” (the denial of the Neutrality Thesis, **C**).

Part iv: Problems for Mulhall’s Argument Against the Neutrality Thesis

De Mesel (2015) criticises Mulhall’s argument on the grounds that that Mulhall’s conclusion is not interesting, since it is (De Mesel thinks) not controversial: ‘I am quite sure that most contemporary meta-ethicists would object to the idea that there is a sharp

distinction between normative ethics and meta-ethics', he writes, citing Sayre-McCord (2012) to support this claim.

De Mesel's notion of a 'sharp distinction' between normative ethics and meta-ethics, however, is unhelpfully vague, since it does not help us differentiate between opponents of the Neutrality Thesis, like Mulhall, on the one hand, and advocates of the Neutrality Thesis, like Johnston and Schroeder, on the other. Schroeder, for example, agrees there is no sharp distinction between meta-ethics and normative ethics, since he allows that they often intertwine, but he still defends the Neutrality Thesis, since he thinks some meta-ethical positions do not commit those who accept them to any normative ethical claims (2017: 685). Schroeder would hence deny the more radical view, defended by Mulhall and the present paper, that there cannot be an ethically neutral meta-ethics. De Mesel hence gives no good reason to dismiss Mulhall's argument or, for that matter, the argument I will give in the next part of this Section.

However, while De Mesel's criticism of Mulhall is off the mark, I believe there are two serious problems with Mulhall's argument.

The first is that Mulhall does not do enough to defend **Q1**, the idea that a meta-ethical position must satisfy MacIntyre's Requirement. For Mulhall, would-be meta-ethicists who do not satisfy MacIntyre's Requirement misfire, and fail to 'adequately [characterise] the subject-matter of moral philosophy... [on this view] one's notion of what counts as ethical will become capacious to the point of emptiness.' (Mulhall 2002: 303). Mulhall, like MacIntyre, clearly finds it plausible that there *is* a distinction between moral language and manipulative language, and that any (attempt at a) meta-ethical account which does not respect this supposed distinction inevitably misfires. But this is a controversial claim, and Mulhall's assurance that it is "more or less plausible and attractive" (2002: 301) will do little to assuage its critics. We saw already in the last part of this Section, for example, that the likes of Stevenson (1945) are committed to meta-ethical views which violate MacIntyre's Requirement. From what Mulhall tells us, it is clear that meta-ethicists must choose between MacIntyre's Requirement and Stevenson's meta-ethics, but it is not clear

why the former is necessarily for choice. Why could the likes of Stevenson not simply take his arguments for his emotivist meta-ethics to be a *reductio* against MacIntyre's Requirement?

The fact that Mulhall does not do enough to defend **Q1** is already a serious problem. It would hence benefit opponents of the Neutrality Thesis to have an argument which works independently of whether **Q1** is true, at least unless some comprehensive defence of **Q1** is in the offing. My argument below works independently of whether **Q1** is true, hence removing the need for opponents of the Neutrality Thesis to commit to **Q1**.

The second problem with Mulhall's argument is that I do not think he convincingly explains why **Q2** is true, i.e. why a meta-ethical position which satisfies MacIntyre's Requirement *cannot* be ethically neutral. From his discussion, we get a sense that perhaps accepting MacIntyre's Requirement *might* commit one to the normative ethical claim that manipulation is morally wrong. Mulhall claims that there is a "Kantian resonance" in the language MacIntyre uses to explain the (supposed) distinction between moral and manipulative language, and hence that accepting that distinction "would amount to endorsing... a morally-charged and motivated contrast" (2002: 302).

We can sensibly ask, though, whether using language with a "Kantian resonance" really commits one to any normative ethical claims. Admittedly, it is easy to think of a case of somebody who endorsed MacIntyre's distinction between moral and manipulative language for "morally charged" reasons. For example, somebody who believes manipulation is always a serious moral wrong might insist on that distinction because they believe the distinction helps us respond appropriately to manipulative behaviour. Mulhall's claim, though, is not simply that the distinction *can* be endorsed for "morally charged" reasons, but that it is *always* (at least in part) endorsed for "morally charged" reasons, for if MacIntyre's Requirement could be endorsed without a normative ethical commitment, then **Q2**, the premise Mulhall needs for his argument, would be false (since it would then be possible for a meta-ethical position to accept MacIntyre's Requirement while still being ethically neutral).

Since Mulhall does not give us a convincing reason to accept **Q2**, then, opponents of the Neutrality Thesis need to make clearer why a meta-ethical position which accepts MacIntyre's Requirement *cannot* be ethically neutral (note the emphasis on the "cannot" as opposed to "may not"). My argument against the Neutrality Thesis will also attempt to improve upon Mulhall's in this respect.

In this part and the previous part of the present Section, I introduced and explained "MacIntyre's Requirement" that meta-ethical accounts respect the distinction between moral and manipulative language. We saw that Cavell (1979) and MacIntyre (1981) object to Stevenson's (1945) emotivism, for example, because it violates that distinction, and they find this objectionable on both meta-ethical grounds (because they think this is a conceptual mistake) and normative ethical grounds (because they think undermining this distinction has problematic social consequences). Mulhall then argues that there cannot be an ethically neutral meta-ethics, since (he claims) any adequate meta-ethical account must accept MacIntyre's Requirement, but accepting this Requirement is non-neutral. I then outlined two serious problems with his argument. I shall now offer a new argument against the Neutrality Thesis which, while similar to Mulhall's, avoids the problems I just outlined.

Part v: A New Argument Against the Neutrality Thesis

Having explained the controversy surrounding the distinction between moral and manipulative language, I can now present my argument.

If we are to make some meta-ethical claim, I argued in part ii) of this Section, we must have something to say about the distinction between moral and non-moral language. The claim that 'moral judgments do not express beliefs', for example, is empty if we do not have something plausible to say about which judgments are, and are not, moral ones. What we say about the difference between moral and non-moral language will either satisfy MacIntyre's Requirement or it will not. Whether or not one's account of the difference

between moral and non-moral language satisfies MacIntyre's Requirement is a binary matter.⁵⁰ In the case that a meta-ethical theorist does not explicitly address the distinction between moral and manipulative language, we can still work out whether their account satisfies MacIntyre's Requirement, since as long as the meta-ethical account in question has something substantive to say about the difference between moral and non-moral language, this can be used as a basis for determining the account's compatibility with the Requirement.

This is summed up in the first premise of my argument:

P1) A meta-ethical position either satisfies *MacIntyre's Requirement*, or it does not.

By the phrase 'meta-ethical position', I simply mean a bundle of one or more meta-ethical claims. If one has made one or more meta-ethical claims, I am happy to grant that one has a meta-ethical position. My point is that to make a (substantive) meta-ethical claim, we must have something to say about the difference between moral and non-moral language, and this either satisfies MacIntyre's Requirement or it does not.

If this first premise is true (and I believe it is hard to doubt), then we can reach the conclusion that there cannot be a neutral meta-ethical position by adding the following two premises:⁵¹

P2) If a meta-ethical position satisfies *MacIntyre's Requirement*, it is not ethically neutral.

⁵⁰ For some accounts, especially minimal ones, it might be difficult to tell whether they conform to the requirement. However, it being difficult to tell whether an account respects the distinction does not imply that there is no truth of the matter about whether the account respects the requirement or not.

⁵¹ There are other arguments which could give us the same conclusion, one of which is made in Mulhall (2002). I will discuss my reasons for preferring my argument over Mulhall's argument in part v), below.

P3) If a meta-ethical position does not satisfy *MacIntyre's Requirement*, it is not ethically neutral.

Jointly, **P1**, **P2** and **P3** entail:

C) There cannot be an ethically neutral meta-ethical position.

This conclusion would show that, contrary to the Neutrality Thesis, an ethically neutral meta-ethics is not possible. Would-be neutral meta-ethicists, I argue, either satisfy MacIntyre's Requirement, or they do not. Either way, they cannot maintain ethical neutrality.

Note how this version of the argument avoids the first problem for Mulhall's argument that I outlined in the previous part of this Section. Since I believe that neutrality cannot be maintained whether or not meta-ethicists accept MacIntyre's Requirement, my argument (unlike Mulhall's) does not depend on MacIntyre being *right* to insist on his Requirement.

Of course, **P2** and **P3** require arguments to support them, since they are not obviously true. In the rest of this Section, I try to show that they are both true. (By arguing for **P2**, I also solve the second problem for Mulhall's argument discussed in the last part, namely that he does not satisfactorily explain why accepting MacIntyre's Requirement cannot be neutral).

Part v, Sub-Section i: Why Satisfying MacIntyre's Requirement is not Ethically Neutral

As I mentioned in parts i) and ii) of this Section, one's meta-ethical position is ethically neutral if and only if there are no normative ethical claims with which it is inconsistent.

What I aim to show, then, is that any meta-ethical position which satisfies MacIntyre's Requirement is inconsistent with at least one normative ethical claim.

Since satisfying MacIntyre's Requirement means conforming to the distinction between moral and manipulative language, and MacIntyre insists on this Requirement partly due to a normative ethical concern about manipulative behaviour and social relations (see Section 3, above), we might wonder whether satisfying MacIntyre's Requirement commits us to a normative ethical view which says that manipulation is morally wrong, or something in that territory. This is exactly what Mulhall alleges,⁵² and I think he is basically correct, but my reasoning here is different to his (see part v), below, for further explanation).

Consider the normative ethical view *anti-Kantianism*:

*It is morally right to manipulate others, and to allow others to manipulate you.*⁵³

I will argue that satisfying MacIntyre's Requirement commits one to denying anti-Kantianism. First, though, let me clarify the anti-Kantian view. It is not just that anti-Kantians say manipulation is just *one* of the morally right things to do. Rather, it sees manipulateness as *the* central right-making feature of right actions and practices,⁵⁴ just as Divine Command Theorists think the central right-making feature of right actions and practices is that they are commanded by God.

I stipulate that the anti-Kantian, like the Kantian, understands by 'manipulation' the treating of persons as mere⁵⁵ means (cf. part iii), above). The anti-Kantian differs from the

⁵² Mulhall (2002: 302).

⁵³ I call it anti-Kantianism because it is the opposite of Kantian first-order ethics, since Kant urges us to treat humanity, both in oneself and in other persons, always as ends in themselves and not as mere means. So Kant urges *against* manipulative relations between any human beings, no matter whether the self is manipulator or manipulated. See Kant (2012 [1785]).

⁵⁴ Anti-Kantianism is thus similar to the view that 'justice is the rule of the stronger', defended by Thrasymachus in Chapter 1 of Plato (2007 [381BC]).

⁵⁵ I do not intend "mere" to be a normative term here: for example, I am not using the word "mere" to suggest that treating somebody as a means must be disrespectful, insulting, or otherwise wrong. Rather, I

Kantian not in their understanding of what manipulation is, but in their view of its moral status. Unlike the Kantian, the anti-Kantian believes manipulative behaviour to be morally good.

On this understanding of ‘manipulation’, it is *not* part of the definition of manipulation that it is morally wrong. If that *were* part of the definition of the concept, then anti-Kantianism would not be coherent, so it would not be a genuine normative ethical view. I will consider another way of challenging the idea that anti-Kantianism is a normative ethical view in part v).

While anti-Kantianism is not plausible, this does not matter for my argument. I just need to prove that making one’s meta-ethics satisfy MacIntyre’s Requirement is inconsistent with at least one normative ethical claim, and my point here is just that anti-Kantianism appears to be a coherent normative ethical claim: like many other normative ethical views, it picks out a certain kind of action (think of e.g. ‘maximising happiness’, ‘doing what God commands’, etc.) and says that kind of action and behaviour is morally good. Unless we find special reason to think otherwise, then, we are safe in assuming that anti-Kantianism is a normative ethical view (if an implausible one).

Faced with a statement such as ‘you should give me money’, MacIntyre, and those who conform to his Requirement, would note an ambiguity between at least two possible meanings. As we saw in part iii), this could be a manipulative statement, being used to convey something like ‘you should give me money, or else’. Alternatively, perhaps the speaker means to call attention to a genuine moral reason for the addressee to give them money – in which case, MacIntyre will say, it is not a manipulative but a moral statement.

believe that it is possible to treat somebody as a means without treating them as a *mere* means. I may treat the dinner lady in my office cafeteria as a means to getting my lunch, but so long as I pay regard to her agency and treat her with respect, I may not be treating her as a *mere* means (since I would also be treating her as an end at the same time). A classic example of treating a person as a *mere* means is the way a slave-owner treats his slave (i.e. with no regard for the slave’s own agency or needs).

This verdict would be disputed by the anti-Kantian. For the anti-Kantian, the manipulative version of ‘you should give me money’ is both manipulative *and* ethical at the same time (and indeed, is ethical *because* it is manipulative). What is the source of this disagreement?

The natural thought behind MacIntyre’s view is that there is a great difference between *forcefully expressing one’s will for somebody to do something* and *identifying a moral reason for somebody to do something*. Conventionally, we would say that somebody who issues a threat to you, or otherwise manipulates you, does not thereby give you a *moral* reason to do what they want you to do, even if they do thereby give you a *practical* reason to do it. This widespread intuition is appealed to, for example, in discussion about ‘Wrong Kind of Reason’ (WKR) cases.⁵⁶

Crucially, however, anti-Kantians must reject this conventional view. On their view, to forcefully express one’s will for somebody to, say, give them money – e.g. by issuing a threat – *is also* to give them a moral reason to give you money. This is because, according to anti-Kantianism, manipulative utterances, behaviour and actions are sources of moral reasons, since (the view says) it is morally right for others to manipulate you (treat you as a mere means).

To use a parallel, it would clearly be inconsistent with Divine Command Theory to deny that it being God’s will that you do X provides moral reason(s) for you to do X. For a Divine Command Theorist, God commanding you to do X always, necessarily, means you have moral reason to do X. In the same way, for anti-Kantianism, that a person A is trying to manipulate you to do X always, necessarily, means you have a moral reason to do X. Hence, it is inconsistent with anti-Kantianism to deny that somebody trying to manipulate you to do X provides you with moral reason(s) to do X.⁵⁷ Yet, due to how MacIntyre’s

⁵⁶ See Hieronymi (2005: 437-457). Hieronymi’s basic thought is that there is a difference in type between practical reasons to do certain things, e.g. when one has a reason to drink a cup of mud because one is being offered a large sum of money to do so, and moral reasons to do certain things, e.g. when doing a certain action would be intrinsically rewarding or fulfilling (unlike drinking a cup of mud). This thought is in the spirit of MacIntyre’s distinction between moral and manipulative language. The term “Wrong Kind of Reason” was initially introduced by Rabinowicz et. al. (2004).

⁵⁷ In this context, I mean by ‘reason’ what Parfit (1984: 153) calls ‘objective’ reasons. These are reasons that one can have independently of your being aware of them. Perhaps you are not aware that God commands that

Requirement insists on a distinction between manipulative uses of “should” (and similar terms) and ethical uses, it *does* deny this, and it is therefore incompatible with anti-Kantianism.

To see this, consider again how MacIntyre treats statements like ‘you should give me money’. As far as he is concerned, this statement could mean ‘you should give me money, or else’ (and is manipulative) *or* it is offered as a sincere attempt to draw attention to a moral reason for giving the speaker money (and it is moral).⁵⁸ The underlying assumption here is that manipulative utterances⁵⁹ are not sources of moral reasons. Otherwise, why would one distinguish so sharply between a use of ‘you should give me money’ which could relate and draw attention to moral reasons, and a manipulative use? The anti-Kantian can, therefore, complain that MacIntyre’s view stacks the deck against her. MacIntyre’s separation of two kinds of ‘should’-statement, manipulative and moral, presupposes that manipulative considerations are not sources of moral reasons. But this is exactly what anti-Kantianism denies, and we hence have an incompatibility here.

So Mulhall’s suspicion⁶⁰ turns out to be well-founded: MacIntyre’s criticism of Stevenson, and the Requirement he places on accounts of moral language, are not morally neutral. What I have done here is to draw out, in more detail than Mulhall, exactly why this is so. Rather than appealing to Mulhall’s idea that accepting MacIntyre’s Requirement involves using language with “Kantian resonance”, I have instead argued that accepting MacIntyre’s Requirement conflicts with the normative ethical view I called ‘anti-Kantianism’. Hence, if a meta-ethical position satisfies MacIntyre’s Requirement, it is not morally neutral (so **P2** of my argument is true).

you do X, but according to Divine Command Theory, you still have objective reason to do X, even if you are not be aware of this.

⁵⁸ There are other potential meanings the sentence could have. For instance, the “should” could be a *prudential* “should”. But the relevant point here is that, for MacIntyre, the sentence could have either of the two meanings I initially mentioned.

⁵⁹ i.e. those which are typically used to treat others as mere means, e.g. ‘give me money, or else!’

⁶⁰ See Mulhall (2002: 301-3).

Part v, Sub-Section ii: Why Not Satisfying MacIntyre's Requirement is Not Ethically Neutral Either

I have just argued that MacIntyre's Requirement is tied to a specific normative ethical perspective, one which, broadly speaking, disapproves of manipulative behaviour and is thus inconsistent with anti-Kantianism, which approves of manipulative behaviour. We might wonder, then, whether failing to satisfy MacIntyre's Requirement is inconsistent with the anti-manipulative ethic that informs MacIntyre's criticism of Stevenson. Specifically, I will argue that failing to satisfy MacIntyre's Requirement is inconsistent with the *Stringent Anti-Manipulative View*:

It is never morally permissible to hinder the recognition, or criticism, of manipulation.

A meta-ethical position which fails to satisfy MacIntyre's Requirement, like the Stevensonian view, does not distinguish between moral and manipulative language. As we saw in Section III, this is morally controversial on the grounds that it hinders our ability to recognise, and critically respond to, acts of (attempted) manipulation. By equating 'X is morally good' with 'I approve of X, do so as well!', Stevenson arguably⁶¹ blurs the boundary between moral and manipulative language, reducing moral language to the mere expression of the speaker's will for others to approve of what she approves of. By erasing this apparent boundary, MacIntyre would complain, Stevenson's view makes it more difficult to tell whether a given speech-act is aiming for manipulation, or the drawing of attention to an (apparent) moral reason. (Stevenson and his followers, meanwhile, would likely deny that these are two different things). For those who are morally concerned about manipulation, this is a morally regrettable consequence of Stevenson's view, and the 'Stringent Anti-Manipulative View' is one example of a moral view which would rule against failing to satisfy MacIntyre's Requirement.⁶² Hence, **P3** of my argument against

⁶¹ The reason I say "arguably" is that, unlike Mulhall, I do not need to take a stance on whether meta-ethicists must accept MacIntyre's Requirement in order for my argument against the Neutrality Thesis to work (see the opening paragraphs of part v), above).

⁶² As with anti-Kantianism, the Stringent Anti-Manipulative View is not very plausible. (It would be more plausible to say that it is at least sometimes morally impermissible to hinder our ability to recognise, and critically respond to, manipulation). But the fact that the Stringent Anti-Manipulative View is not very

the Neutrality Thesis is true: a meta-ethical position which does not accept MacIntyre's Requirement (to respect the distinction between moral and manipulative language) is not ethically neutral.

Imagine the would-be neutral meta-ethicist who refuses to make their meta-ethics satisfy MacIntyre's Requirement tries to protest that their neutrality is still intact. They may say that, in giving their meta-ethical claims, they pass no judgment on manipulation, or any associated moral issues. However, while they may pass no judgment on such issues explicitly - on the page - the point is that by conducting their meta-ethics in a certain way, they are doing something that at least one first-order morally view deems morally impermissible (i.e. failing to preserve the distinction between manipulative and moral language). It would hence be futile for them to claim moral neutrality, when their meta-ethical practice is in clear tension with at least one normative ethical view.

If they genuinely believed that it is morally impermissible to undermine the distinction between manipulative and moral language, why would they not make their meta-ethics satisfy MacIntyre's Requirement? Otherwise there would be a clear incompatibility between their meta-ethical conduct and their first-order moral views. While a situation of such incompatibility is, of course, psychologically possible (as is familiar from weak-will cases), the incompatibility between the moral view and the conduct is real enough. A fully rational believer of the Stringent Anti-Manipulative View would make their meta-ethics satisfy MacIntyre's Requirement instead.

This explains why the Neutrality Thesis is plausible, despite being false. The likes of Johnston (1989) who try to give a neutral meta-ethics while endorsing a Stevensonian, broad account of what counts as moral language,⁶³ do not have to *explicitly* endorse any first-order moral claims in the course of expounding their meta-ethical claims and arguments, which can give their readers the impression that their meta-ethical account is

plausible does not matter for my argument: I do not assume that it is true, only that it is a coherent, genuine normative ethical view.

⁶³ See part iii) above for discussion of Stevenson's broad way of delineating moral from non-moral language, and cf. Mulhall's comment at (2002: 301).

itself ethically neutral. This is not the case, however. The drawing of certain key distinctions, such as that between what does and does not count as moral language, is itself subject to moral objections and concerns,⁶⁴ as we have seen with the example of the distinction between moral and manipulative language. The normative ethical commitment in Johnston's view is *implicit* rather than explicit, but as our investigation of MacIntyre's Requirement has shown, it is real enough.

Again, consider how futile it would be to claim moral neutrality on some issue if it could be shown that our statements (or indeed other actions or behaviour) regarding that issue were deemed morally impermissible by some normative ethical view. While we might try to protest our neutrality in such a case, this would be entirely unconvincing. If I have just stolen your car, it would be absurd for me to claim to be neutral on the normative ethical issue of whether stealing cars is permissible. A similar point applies to the meta-ethicist who does not meet MacIntyre's Requirement who professes ethical neutrality. Just as there is a clear conflict between my conduct in stealing the car and the claim that stealing cars is impermissible, so too there is a conflict between failing to meet MacIntyre's Requirement and the Stringent Anti-Manipulative normative ethical view.

My argument in full, then, is as follows:

- P1)** A meta-ethical position either satisfies *MacIntyre's Requirement*, or it does not.
- P2)** If a meta-ethical position satisfies *MacIntyre's Requirement*, it is not ethically neutral.
- P3)** If a meta-ethical position does not satisfy *MacIntyre's Requirement*, it is not ethically neutral.
- C)** There cannot be an ethically neutral meta-ethical position.

To briefly recap: if we are to count as making anything that deserves to be called a 'meta-ethical' claim, we must have something to say about what differentiates moral from non-

⁶⁴ Cf. Mulhall (2002: 302).

moral language. What we say about this, though, will reveal our view on the relationship between moral language and manipulative language. Do we side with MacIntyre, and see these as utterly distinct and separate? Or do we, like Stevenson, not account for any fundamental difference between the two? Without passing judgment on what the right answer to this question is, I have argued that either way, our meta-ethical position will not and cannot be ethically neutral.

If we satisfy MacIntyre's Requirement, we are assuming that manipulative utterances are not sources of moral reasons, something denied by the normative ethical view I called anti-Kantianism. If we do not satisfy MacIntyre's Requirement, we are displaying a certain tolerance towards the risk of undermining our ability to recognise, and critically respond to, manipulative utterances and behaviour, and this is deemed morally impermissible by the Stringent Anti-Manipulative View.

Having defended my argument against the Neutrality Thesis, I will now set out why it is an improvement upon the previous argument against the Neutrality Thesis offered by Mulhall (2002).

First, in part iv) I argued that Mulhall's argument depends on the claim that all meta-ethicists must accept MacIntyre's Requirement to respect the distinction between moral and manipulative language, a claim that Mulhall does not argue for. The argument I have developed does not depend on this claim.

I hence remain officially neutral on whether meta-ethicists should or must accept "MacIntyre's Requirement" to respect the distinction between moral and manipulative language. However, it is important to note that in case it is true that meta-ethicists must accept MacIntyre's Requirement (i.e. Mulhall's **Q1** is correct), then Mulhall's argument against the Neutrality Thesis may be sufficient on its own. My argument gives opponents of the Neutrality Thesis the option of not committing to Mulhall's **Q1** (the claim that meta-ethicists must respect MacIntyre's Requirement), but I am not attempting to determine the

truth or falsity of **Q1** here, so how much of an advantage this is for my argument is an open question.

A further advantage of my argument is that it provides a deeper explanation than Mulhall's of why both satisfying and not satisfying MacIntyre's Requirement in one's meta-ethics compromises one's neutrality.

What my argument against the Neutrality Thesis achieves, in the context of this thesis, is that it casts doubt on Johnston's (1989) view that his Emotivist meta-ethical position is, unlike realist positions, ethically neutral (see Section 2, above). Since this is one of the main advantages Johnston claims for his position, this is an important result which makes Johnston's position less attractive than it might seem to be at first glance. Moreover, what we have found makes it seem less troubling if realist positions do involve normative ethical commitments⁶⁵ – after all, if *all* meta-ethical positions carry such commitments, this cannot be a unique fault with realism.

Moreover, if every meta-ethical position does in fact conflict with at least some normative ethical claims, this raises the question of which normative ethical claims conflict with Johnston's purportedly-neutral "Emotivist" position. As it happens, answering this question will help me expose some further problems for Johnston's view, and make the path to rejecting the Fact/Value Dichotomy, which lies at the core of Johnston's position, clearer. I hence turn to these matters in the following and final Section of this Chapter.

⁶⁵ I granted Johnston that they do so in Chapter 1, Section 2, part v).

Chapter 1, Section 4: The Ethical Non-Neutrality of Johnston's Wittgensteinian Approach to Moral Philosophy

Part i: The Rationale of this Section

In this Section, I am going to argue that Johnston's meta-ethical position is not ethically neutral, because it is incompatible with at least one normative ethical view. Specifically, I will argue that Johnston's central thesis, that there is an irreducible logical gap between descriptive language and evaluative language, is incompatible with at least some Classical Teleological normative ethical views. The broader implications of this for my thesis are twofold. First, this brings out a hitherto unappreciated cost of adopting Johnston's position, namely that it blocks one from accepting some Classical Teleological normative ethical views. Secondly, it reveals how Johnston must commit to a Separabilist "splitting strategy" with regards to descriptive and evaluative language, the viability of which will be explored in the remaining Chapters of this thesis.

In previous sections of this Chapter, I have argued that Johnston's (1989) Wittgensteinian approach to moral philosophy commits to the idea that we can conduct a second-order, descriptive sort of moral philosophy which is entirely neutral on first-order, evaluative ethical matters. Johnston maintains that his own descriptive account of the grammar of our moral language is entirely neutral in just that way. However, I then argued – taking my cue from Mulhall (2002) - that it is impossible to give an entirely ethically neutral second-order account of our moral concepts.

If this is all correct, then we must understand second-order, descriptive approaches to moral philosophy and accounts of our moral concepts ("meta-ethics" for short) differently to Johnston. We can no longer, with Johnston, view second-order accounts as *merely* descriptive, since they will inevitably have some first-order, evaluative commitments built into them. This naturally raises questions about the first-order evaluative commitments that

may be built into Johnston's meta-ethics. Are there any (unacknowledged) first-order commitments present in Johnston's position? (If my argument above is correct, there must be at least some present). If so, what are they?

Here is my rationale for pursuing these questions. If Mulhall and I are correct that there are first-order commitments hidden under the surface of Johnston's position, we need to be explicit about what these are if we are to fully understand Johnston's position. Secondly, one apparent attractive feature of Johnston's position is that it seemed it could not be controverted on first-order grounds. However, if Mulhall and I are correct, this is not the case. We need to set out the first-order commitments of Johnston's position so we can see why that position may be controversial for first-order reasons. Plainly, understanding all the features, advantages and disadvantages of Johnston's view is also important for weighing up the virtues and vices of Johnston's approach to meta-ethics against potential alternative approaches, such as MacIntyre's.

The central thesis of Johnston's book is that there is an irreducible logical distinction between descriptive language, such as empirical factual statements, and evaluative language, such as moral claims (see Ch. 1 Section 2, above). On the face of it, this appears to be a second-order, descriptive claim. After all, it appears to be a claim about our language and how it works: in Johnston's terminology, it is a "grammatical" claim. It does not appear to commit Johnston to any evaluative claims, and it does not include any of the central evaluative words which usually indicate that a claim is evaluative, such as "good", "ought", "should" etc.

Nevertheless, I believe that this aspect of Johnston's position *does* commit him to certain first-order views. In the rest of this Section, I will explain why.

When defending Mulhall's argument in Ch. 1, Section 3, I appealed to the following criterion to determine whether a claim (such as an apparently second-order claim) is ethically neutral: a second-order position is ethically neutral if and only if there is no first-

order ethical view with which it is inconsistent. I now want to apply the same criterion to Johnston's claim that there is an irreducible logical distinction between descriptive language and evaluative language, and argue that there *is* at least one first-order view which is inconsistent with that claim. I will argue that Johnston's claim is incompatible with at least some Classical teleological normative ethical claims. (I understand Classical teleological ethical views to be a family of similar first-order ethical views which includes Aristotelianism and Thomism; for more detail see below and MacIntyre 1981).

To see why these views are inconsistent, we will need to grasp some further details both of Johnston's claim and of Classical teleological ethical views.

Part ii: Johnston on Inferences from Descriptive Premises to Evaluative Conclusions

First, recall how Johnston's claim leads him to consider deductive inference from descriptive premises to evaluative conclusions to be outright impossible. He writes that "there is always a logical gap between the discoveries of science and the judgements and actions one may derive from them" (1989: 154). Johnston is adamant that the attempt "to claim that our [ethical] concepts and our practices reflect reality or are justified by it" is "incoherent" (1989: 85), and therefore thinks that the project of trying to support an ethical viewpoint with appeal to factual considerations is futile. Johnston adds that "no ethic can be supported on logical or conceptual grounds" (1989: 209), which confirms that Johnston considers the attempt to draw *deductive* inferences from descriptive premises to evaluative conclusions, in particular, doomed to failure (see Ch. 1 Section 1, above).⁶⁶

It is important that it is not *merely* that Johnston thinks such inferences are *illicit*, but rather that he thinks that *attempting* to make such inferences necessarily involves labouring under

⁶⁶ One can, of course, attempt to infer evaluative conclusions from descriptive premises non-deductively, but that is not my focus in this thesis.

a conceptual confusion, which results only in nonsense. *If* we grant Johnston that there is an irreducible logical distinction between descriptive language and evaluative language, I think it follows that there can be no inferences from descriptive premises to evaluative conclusions, for the following reason.

If Johnston were correct that descriptive and evaluative language are logically separate, then as Hume complains, it would be utterly mysterious why we could bring in a new, evaluative term ("ought") into our conclusion, when no such thing appeared in the premises (Hume 1985 [1739], Bk. 3, sect. 1, part 1). From this Humean point of view, since descriptive language and evaluative language are simply logically separate, the whole enterprise of trying to make deductive inferences from descriptive premises to evaluative conclusions is no better than trying to magic a rabbit out of a hat. There is, according to this view, simply no evaluative material to be found in bare descriptive premises from which one can muster any evaluative conclusions. Thus, it is natural enough for Johnston to move from the idea that there is a "logical gap" between descriptive and evaluative language to the conclusion that there can be no deductive inferences from descriptive premises to evaluative conclusions.

One possible reason favouring this Johnston-style scepticism about inferences from "is" premises to "ought" conclusions is the principle that, in a logically valid argument, "nothing may appear in the conclusion which does not appear in the premises" (for discussion see MacIntyre 1959, 1981: Ch.5). If this is correct, then we will have to accept that, *if* descriptive language and evaluative language are entirely logically distinct and separate, one cannot infer from descriptive premises to evaluative conclusions: since, in case this is true, there would be no evaluative material in any descriptive premises, and we could not "magically" introduce anything evaluative into the conclusion if there was nothing evaluative in the premises.

MacIntyre, however, protests that the principle that, in a valid argument, nothing may appear in the conclusion which does not appear in the premises is "bogus", and in a certain sense he is correct. As Anscombe (1958) points out, there are perfectly legitimate

inferences in which terms appear in the conclusion which do not appear in the premises. Here is an adapted version of her example:

- 1) You ordered £20 worth of potatoes from your grocer.

- 2) Your grocer left £20 worth of potatoes at your property along with a bill for £20.

- C) You owe your grocer £20.

As Anscombe points out, a new term ("owes") appears in the conclusion of this argument which does not appear in the premises. Yet it would surely be wrong-headed to deny that this is a perfectly good inference! So Anscombe and MacIntyre think that there can be logically valid inferences in which a term appears in the conclusion which does not appear in the premises.

Although I think there is something right about this,⁶⁷ Anscombe and MacIntyre seem to overlook the possibility that what is going on in such cases is that there are suppressed premises, which explain the validity of the inference. Consider again Anscombe's grocer example. Is there not a suppressed premise, as follows?

- 1) You ordered £20 worth of potatoes from your grocer.

- 2) Your grocer left £20 worth of potatoes at your property along with a bill for £20.

⁶⁷ I return to this issue in Chapter 2, Section 1 and Chapter 3, Section 3, below. My conclusion is ultimately sympathetic to Anscombe and MacIntyre's position, but I think these cases require slightly different treatment to what they offer.

3) If you order £X worth of a good from a supplier, and that supplier delivers those goods to you and leaves a bill for £X, you owe that supplier £X.

C) You owe your grocer £20.

Adding the suppressed premise, 3, makes this a logically valid argument by the usual definition, i.e. that its conclusion cannot be false if all the premises are true. (The first argument is hence not logically valid, even if it seems like an inference that “works” for practical purposes in everyday conversations). It seems then that, after all, we *cannot* have new terms appearing in the conclusions of logically valid arguments which did not appear in the premises. There is a further discussion to be had about these types of arguments, which I will call “Anscombe/MacIntyre” arguments,⁶⁸ but for the purposes of this Section, let us note that Johnston at least appears to be correct that *if there is* a “logical gap” between descriptive language and evaluative language, the whole business of deductively inferring evaluative conclusions from descriptive premises is a non-starter.

My strategy for showing how Johnston's position is not ethically neutral involves pointing out how the views of Johnston we have just recapped rule out at least one first-order ethical view. It is the commitment of Johnston's idea that there cannot be deductive inferences from descriptive premises to evaluative conclusions which I think is incompatible with at least some versions of one kind of normative ethical view: namely, Classical teleological normative ethical views. In part iv, I set out why I perceive there to be an incompatibility here, but first we must go over some relevant core features of Classical teleological ethical views.

⁶⁸ See Chapter 2, Section 1 and Chapter 3, Section 3, below.

Part iii: Classical Teleological Normative Ethical Views

I use the label "Classical teleological ethical views" to refer to a family of related ethical views grouped together by MacIntyre (1981: 62). Members of this family include Aristotelianism, Augustinian Christian ethics, and Thomism (MacIntyre 1981: Ch. 5, MacIntyre 1986). MacIntyre claims that these views share a basic common structure, namely the one "Aristotle analyzed in *Nicomachean Ethics*" (1981: 62). I will now attempt to elucidate this common structure, so we can understand why views in this family are incompatible with Johnston's supposedly-neutral meta-ethics.

Before I begin this elucidation, I want to mention a caveat. Normative ethical theories such as Aristotelianism and Augustinianism are, themselves, not monoliths, and there is significant diversity and variation even within the group of, say, Aristotelian ethicists (see e.g. Kenny 2001). I neither want or need to ignore this diversity, since, first of all, I aim only to elucidate some *common features* of Aristotelian and classical teleological views, *not* to give a comprehensive analysis of exactly which claims a true Aristotelian must subscribe to. Secondly, since I only need to find *one* normative ethical view which is incompatible with Johnston's view to prove my case, the potential existence of atypical Classical teleological views, which may not be well-captured by the following rough elucidation, does not cause trouble for my argument against Johnston.

It is also important to note that there are significant differences between sub-traditions of the Classical teleological family (see e.g. MacIntyre 1988: Ch. IX). For the same reason mentioned in the above paragraph, though, this is not a problem for my argument, as long as I can pinpoint *some* Classical teleological views which are inconsistent with Johnston's view.

What common features unite Classical teleological views? According to MacIntyre, at the broadest level these views all revolve around a contrast between "man-as-he-happens-to-be" and "man-as-he-could-be-if-he-realized-his-telos" (1981: 62-3). Classical teleologists

view ethics as the science which enables human beings to "transition from the former to the latter". Classical teleologists, then, aim to provide us with three things: i) an account of human nature ("man-as-he-happens-to-be"), which informs us about what we are and what we can be, ii) an account of our *telos*, or ultimate end, which is conceived of as a state of ethical perfection, and iii) an account of how we can move from our present state towards the *telos*. In the Classical teleological ethical tradition, the concept of virtue is central. One of the main reasons for this is that Classical teleologists claim that the possession and exercise of the virtues are necessary conditions for reaching one's *telos* and that exercise of the virtues partially constitutes our *telos* (see e.g. MacIntyre 1981: Ch. 14).⁶⁹

Classical teleologists hence understand the moral life as the journey from our starting point towards the *telos*; a journey in which the virtues are essential in guiding us towards, and helping us realise, our *telos*. But how do these theorists conceive of the *telos* and, indeed, the "starting point", i.e. "man-as-he-happens-to-be"? This is a point on which Classical teleologists are divided,⁷⁰ at least when it comes to the details, but the following broad ideas are commonplace. Human beings exist in a natural state, but a state in which our natures are imperfect in some way. To move further towards the *telos* is for one's *nature* to improve; "virtue in general is nature perfected", as MacIntyre (1988: 148) explains Cicero's view. This state of perfected nature, the *telos*, is understood as a state of *harmony with nature* or sometimes alternatively as *accordance* with nature. As we move towards the *telos*, then, the extent to which we are in harmony with nature grows. An ethical agent, on this picture, irons out creases where harmony between herself and nature is presently lacking, thereby moving closer to the *telos*. As the Stoic Diogenes Laertius puts it, "to live in conformity with nature means to live a life of virtue, since it is to nature that virtue leads" (Hicks 1910).

How do Classical teleologists understand this crucial notion of "harmony" with nature?

Typically, Classical teleologists understand the process of living in greater harmony with

⁶⁹ Cf. Aristotle's claim that *eudaimonia* (roughly: happiness) is "activity in accordance with virtue"; cited in McDowell (2009: 25).

⁷⁰ The introduction of the notion of original sin by Christian thinkers like Augustine complicates matters. See MacIntyre (1988) for explanation and discussion.

nature to involve embracing and developing one's natural skills and potentials *qua* human being. At this point it will help to make these abstract ideas more concrete by looking at an example of a Classical teleological ethical view applied to a particular ethical issue.

Consider a hedonist, Harry, who believes that the most ethically justified way for us to live is for us to maximise the amount and intensity of sensual pleasures that we enjoy. So, for instance, Harry believes we should spend most of our time pursuing and enjoying sensual pleasures relating to food, sex etc., and lives accordingly.

To understand Classical teleological ethical systems better, I want to introduce another character, Anna the Aristotelian. Imagine Anna knows of Harry's hedonistic beliefs and lifestyle. Anna might try to convince Harry to alter his hedonistic lifestyle with the following reasoning:

"You are right that appreciating sensual pleasures relating to food, sex etc. is an important part of human life, and of course there is a proper place for these things in the good life. After all, human beings share with most animals in their sensuality and capacity to appreciate sensual pleasures, so a complete and fulfilled human life involves enjoying such pleasures..."

However, a human life dedicated ONLY to pursuit of sensual pleasures is not a good human life. A life dedicated only to pursuit of sensual pleasures is a life fitting for an irrational animal, not a rational human being. As rational beings, humans have the capacity to develop intellectual virtues and appreciate the pleasures involved with contemplation and other intellectual activities. A human being who pursues and appreciates both sensual and intellectual pleasures is living more in harmony with his own nature, and therefore in a more virtuous way, than a human being who only pursues and appreciates sensual pleasures. The reason is that a person who makes room for the intellectual virtues and pleasures is thereby living in a way that fits with his nature as a rational being, whereas a Hedonist, who makes no room for such things, does not – and it

is ethically right to live in the way that best fits with our natures.”

Notice how this reasoning fits with the teleological structure to which MacIntyre calls attention. Our Aristotelian here gives voice to certain views about human nature, “man-as-he-happens-to-be”, i.e. that human beings are rational as well as sensual creatures. These views inform the Aristotelian’s conception of the *telos*: since virtue is “nature perfected”, the *telos* is the state at which the natural capacities of human beings (e.g. rationality) are developed as far as possible and used in the fullest possible way. A hedonistic lifestyle which only prioritises hedonistic pleasures is judged ethically inferior on the basis that it is not suitable for reaching the *telos*, since it neglects the development of intellectual virtues and pleasures.

Hopefully this example makes the abstract idea of “living in harmony with nature” more concrete. According to Aristotelians, since human beings are rational creatures, a life lived in harmony with nature must make room for the activities and virtues which connect with our rational natures, e.g. contemplation, prudence, etc. I explored the classical teleological reasoning behind such claims through the imagined debate between the hedonist and the Aristotelian, above.

Part iv: The Incompatibility of Johnston’s Position and Classical Teleological Views

Having explained the basic structure of Classical teleological views and given an example of how they apply to particular moral controversies, I now want to link our discussion back to the compatibility or otherwise of Johnston’s purportedly neutral meta-ethical position with Classical teleological first-order views.

It is crucial, in this regard, that Classical teleological ethical views *sanction* inferences from descriptive premises to evaluative conclusions. In the Aristotelian criticism of Harry’s hedonism that we considered above, for instance, we see the following inference:

P1) Human beings are rational creatures.

C) Human beings ought to develop the intellectual virtues.

The basic reason that Classical teleologists are inclined to sanction inferences of this form (where, crucially, the premise looks descriptive and the conclusion looks evaluative) is that they *regard nature as a source of ethical authority*. So, for instance, Classical teleologists take the descriptive natural fact that human beings are rational creatures to have direct first-order ethical consequences: for instance, that hedonistic lifestyles such as Harry's are ethically sub-par, owing to their failure to give due weight to developing intellectual virtues. For Classical teleologists, such hedonistic lifestyles are fitting and right only for non-rational animals, not human beings (see e.g. Aquinas, *Summa Theologica* IaIIae, q. 2 a.6).

By contrast, we can imagine those who do not subscribe to Classical teleology questioning whether nature has the kind of ethical authority that Classical teleologists believe it has. Notably, this would mean that such opponents of Classical teleology are sceptical about whether descriptive natural facts are appropriate guides to our conduct, morally speaking. For example, Harry might protest against Anna's criticism along the following lines:

"So what if rationality is part of human nature? I do not think that, just because acting some way is "more natural", it is necessarily ethically superior. So I am unpersuaded by your view that it would be ethically better if I foster the intellectual virtues."

It is not my intention to intervene in the debate between Classical teleology and its critics here, but there is something we need to note about Harry's reasoning: his response involves rejecting a substantive, first-order ethical position. The Classical teleologist

claims that it is ethically right to live and act in a way that fits with one's own nature, while Harry does *not* accept this first-order ethical claim.

The significance of this is that we can now say it appears, at least *prima facie*, that whether one will consider it possible, and legitimate, to infer evaluative conclusions from descriptive premises will depend on the first-order ethical views one accepts. To recap, we found that Classical teleologist ethical views sanction such inferences because they believe nature to be ethically authoritative, and so, for example, they take facts about human nature to have direct⁷¹ normative ethical implications. Those who reject Classical teleology, on the other hand, will not be prepared to accept such inferences. Since they do not accept that nature is ethically authoritative, they will not agree that natural facts, e.g. facts about human nature, have any direct ethical implications. In line with Harry's response to Aristotelian criticism, those who reject Classical teleology will not accept that, for example, just because rationality is part of human nature, it follows that we ought to live in a way that accords with our rationality (e.g. by developing our intellectual virtues).

To state this first version of my argument in a simpler way, call the following inference *CT-I* (Classical Teleologist Inference 1):

⁷¹ The contrast I have in mind between "direct" and "indirect" normative ethical implications can be understood with reference to Johnston's view as set out in Section 2 of this Chapter, above. Recall how Johnston allows that the descriptive claim "causing pollution contributes to rendering the earth uninhabitable" may have implications for how morally ought to act, for instance if we believe that it is wrong to contribute to rendering the earth uninhabitable (1989: 154). But since, on Johnston's view, the descriptive claim on its own lacks any evaluative content, it can only have implications for how we morally ought to act if accompanied by some relevant evaluative claim, like in the pollution example. This would be an example of a descriptive claim having "indirect" normative ethical implications. By contrast, classical teleologists allow that at least some descriptive statements (e.g. those about human nature) *in themselves* have implications for how we morally ought to act. This is what I have in mind with the notion that a descriptive claim has "direct" normative ethical implications.

R1) Human beings are rational creatures.

C) Human beings ought to develop the intellectual virtues.

The thought discussed above is that Classical teleologists endorse inferences such as *CT-1* for first-order ethical reasons: namely that they believe nature to have a kind of ethical authority, and accordingly think that the fact that human beings are rational creatures has direct first-order implications, including that a life with intellectual virtues is, at least *ceteris paribus*, an ethically better life than one without. Johnston, however, owing to his meta-ethical commitment that there cannot be deductive inferences from descriptive premises to evaluative conclusions, would have to deny that this is a legitimate inference. In so doing, however, he reveals that those meta-ethical commitments are incompatible with Classical teleological first-order views. For those meta-ethical commitments rule out inferences like *CT-1*, which (at least most) Classical teleologists must endorse. Hence, Johnston's meta-ethical commitments are inconsistent with at least one first-order ethical view, and hence, by my criterion, are not ethically neutral.

This simple version of the argument, however, does not work. Johnston could reply by making the following point:

CT-1 is not an adequately specified inference. For it to be a valid inference, we would need to bring out some suppressed premises (cf. the discussion of Anscombe's "owes" inference in part ii of this Section), like so:

CT-2:

R1) Human beings are rational creatures.

R2) It is fitting with being a rational creature to develop the intellectual virtues.

R3) If human beings are rational creatures, they ought to do things which are fitting with being a rational creature.

C) Human beings ought to develop the intellectual virtues.

Thus reformulated, the inference is no longer threatening to Johnston's view. For Johnston is perfectly happy to accept that there can be valid arguments for moral conclusions which contain descriptive premises (see Ch.1 Section 2, above); it is only that, for this to be the case, the argument must also include at least one evaluative premise (1989: 154). The evaluative premise allows us to "bridge" the "logical gap" between descriptive and evaluative statements. In *CT-2*, one of the premises certainly is evaluative: R3. As such, Johnston can say, the *CT* inference, once properly unpacked, is perfectly compatible with his own position. The Classical teleologist, he could say, is perfectly free to endorse *CT-2*, but *CT-2* is no threat to his claim that there can be no inferences from exclusively descriptive premises to evaluative conclusions, since *CT-2* has an evaluative premise. So while *CT-1* appeared to cause trouble for Johnston, he now appears to be able to free himself from my critical attack by insisting on reformulating *CT-1* as *CT-2*. Moreover, this appears reasonable to insist upon, since *CT-1* is not a logically valid argument, while *CT-2* is. In other words, in making the *CT* inference logically valid, we also, it seems, make it unthreatening to Johnston's position.

I shall now aim to undermine this line of defence by drawing from Anscombe (1958) and MacIntyre's (1981: 58-61) discussion of the Fact/Value Dichotomy in relation to Classical teleological ethics. To see why Johnston's defence is unsustainable, we first need to consider two different sorts of "suppressed premises" and two different sorts of relation between premises. Consider:

S1) James is a bachelor.

C) James is a man.

Plainly, this is not a logically valid argument – although it is a perfectly decent inference. To make it a logically valid argument we would need to add a “suppressed premise”:

S1) James is a bachelor.

S2) All bachelors are men.

C) James is a man.

Although we have turned our inference into a logically valid argument, in a sense this second formulation says nothing more than the first did. For the “suppressed” S2 simply clarifies the meaning of one of the terms in S1, “bachelor”. Notably, then, in this case there is a special relation between the premises: S2 simply explicates the meaning of S1. This contrasts with the following kind of pair:

T1) Socrates is a man.

C) Socrates is mortal.

T1) Socrates is a man.

T2) All men are mortal.

C) Socrates is mortal.

For here, we would tend to think that T2 is *not* merely unpacking the meaning of T1. It would be odd to think that T2 is an analytic truth, as S2 is, because it does not seem to be part of the meaning of the concept “man” that man is either mortal or, for that matter,

immortal. Although there may never have been immortal men, the idea of an immortal man makes sense, while the idea of a female bachelor does not. By contrast with the “bachelor” case, then, the present case is one where the so-called “suppressed premise” and the original premise are two logically independent claims: one is not a mere explanation of part of the meaning of the other.

The question arises, then, of what sort of relation holds between the premises in *CT-2*:

R1) Human beings are rational creatures.

R2) It is fitting with being a rational creature to develop the intellectual virtues.

R3) If human beings are rational creatures, they ought to do things which are fitting with being a rational creature.

C) Human beings ought to develop the intellectual virtues.

Now, for Johnston, R3 could not be a mere explanation of (part of) the meaning of R1, because R1 is clearly a descriptive claim, while R3 is an evaluative claim. As we know by now, Johnston strictly maintains that there is a logical gap between descriptive and evaluative statements, and would hence regard R1 and R3 as being logically different in kind from one another, and would hence deny that R3 explains the meaning of R1. If Johnston is correct that R3 is not merely unpacking the meaning of R1, then I would grant he would be correct that inferences like *CT-2* pose no threat to his own position.

However, many Classical teleologists would reject Johnston’s view that P1 and P3 are logically separate – and, I hope to show, they do so for first-order ethical reasons, thereby

meeting my aim in this Section of showing how Johnston's position conflicts with Classical teleological ethical views.

To understand this, we must look at MacIntyre's treatment of the Johnstonian idea that there is an irreducible logical gap between descriptive and evaluative language. Against Johnston, MacIntyre would claim that whether, for instance, R3 counts as an explanation of the meaning of R1 (akin to how "all bachelors are male" explains the meaning of "bachelor") or an independent claim (akin to "all men are mortal") depends on whether we treat "human being" as a "functional concept":

... it is helpful to consider another type of counter-example to the "No "ought" conclusions from "is" premises" thesis. From such factual premises as "This watch is grossly inaccurate and irregular in time-keeping" and "This watch is too heavy to carry about comfortably", the evaluative conclusion validly follows that "This is a bad watch". From such factual premises as "He gets a better yield from this crop per acre than any other farmer in the district", "He has the most effective programme of soil renewal yet known" and "His dairy herd wins all the first prizes at the agricultural shows", the evaluative conclusion validly follows that "He is a good farmer".

Both these arguments are valid because of the special character of the concepts of a watch and of a farmer. Such concepts are functional concepts; that is to say, we define both "watch" and "farmer" in terms of the purpose or function which a watch or a farmer are characteristically expected to serve. It follows that the concept of a watch cannot be defined independently of that of a good watch nor the concept of a farmer independently of that of a good farmer; and that the criterion of something's being a watch and the criterion of something's being a good watch – and so also for "farmer" and for all other functional concepts – are not independent of each other. Now clearly both sets of criteria – as is evidenced by the examples given in the last paragraph – are factual. Hence any argument that moves from premises which assert that the appropriate criteria are satisfied to a conclusion which asserts that "that is a good such-and-such", where "such-and-such" picks out an item specified by a functional concept, will be a valid argument which moves from factual premises to an evaluative conclusion. Thus we may safely assert that, if some amended version of the "No "ought" conclusion from "is" premises" principle is to hold good, it must exclude arguments involving functional concepts from its scope. But this suggests strongly that those who have insisted that *all* moral arguments fall within the scope of such a principle may have been doing so, because they took it for granted that *no* moral arguments involve functional

concepts...

Yet moral arguments within the classical, Aristotelian tradition – whether in its Greek or its medieval versions – involve at least one central functional concept, the concept of *man* understood as having an essential nature and an essential purpose or function; and it is when and only when the classical tradition in its integrity has been substantially rejected that moral arguments change their character so that they fall within the scope of the “No “ought” conclusion from “is” premises” principle. That is to say, “man” stands to “good man” as “watch” stands to “good watch” or “farmer” to “good farmer” within the classical tradition. Aristotle takes it as a starting-point for ethical enquiry that the relationship of “man” to “living well” is analogous to that of “harpist” to “playing the harp well” (*Nicomachean Ethics*, 1095a 16) (MacIntyre 1981: 58).

I now want to try to break down this very important passage and its implications for the ethical neutrality (or lack thereof) of Johnston’s position.

Recall that Johnston needs to construe a premise such as “human beings are rational creatures” in *CT-2* as purely descriptive. MacIntyre would beg to differ with this construal, since according to him Aristotelians will understand the concept “human being” (or “man”) as functional, and therefore as having a partly evaluative meaning. As MacIntyre says, “the concept of a watch cannot be defined independently of the concept of a good watch”. The idea here is that it is just *part of what it is*, conceptually speaking, to be a watch that e.g. a watch is for telling the time, and therefore that a good watch tells the time accurately. Along the same lines, MacIntyre thinks that all functional concepts have partly evaluative meaning, and this evaluative meaning is why, on an Aristotelian view, a premise such as “a human being ought to develop the intellectual virtues” is simply explaining the meaning of the concept “human being”; for this is conceived to be an explanation of the functional nature of the concept “human being” in the same way “a good watch tells the time accurately” is, plausibly, an explanation of the meaning of the functional concept “watch”.

The first step in MacIntyre’s argument, then, is to establish that whether or not one will accept that there are counterexamples to the Johnston’s claim (that there can be no

deductive inferences from descriptive premises to evaluative conclusions) depends on whether one accepts that there are descriptive premises which include one or more functional concepts. It seems to me that MacIntyre is clearly right about this. For instance, if “watch” is a functional concept, the inference from “this watch is grossly inaccurate in timekeeping” to “this is a bad watch” is a valid deductive inference. Johnston could try to insist that there is a “suppressed premise” here which is evaluative: “if a watch is grossly inaccurate in timekeeping, it is a bad watch”. However, *if* MacIntyre is right, this so-called “suppressed premise” is already contained within the premise “this watch is grossly inaccurate in timekeeping”, because then if “watch” is a functional concept it is part of the concept “watch” that a watch is bad if it keeps time inaccurately. So, if “watch” is functional, any premises including that concept cannot be “split” into two separate premises, one descriptive and evaluative, in the way Johnston wishes. Rather, we would have to conceive of premises including “watch” and other functional concepts as *both* descriptive and evaluative, and allow that there can be valid deductive inferences from premises which are descriptive (although also evaluative) to evaluative conclusions.^{72 73}

I want to note here that Johnston’s need to “split” certain sentences into separate descriptive and evaluative elements will prove to be crucial when it comes to assessing the Fact/Value Dichotomy (see Chapter 2, below). At present, we seem to be at an impasse: MacIntyre’s idea that concepts like “watch” are somehow both descriptive and evaluative seems *prima facie* plausible, but so does Johnston’s response that sentences featuring concepts like “watch” should be split into separate descriptive and evaluative claims. Note the conditional language I used to explain MacIntyre’s argument in the paragraph above: *if* MacIntyre’s account of concepts such as “watch” is correct, then Johnston’s account should be rejected. But whether MacIntyre’s account of such concepts *is* right after all is, as we shall see in Chapters 2 and 3, difficult to establish.

⁷² To be clear, this view contradicts two views of Johnston; i) that there is a strict logical split between descriptive and evaluative claims and ii) that there can be no valid deductive inferences from descriptive premises to evaluative conclusions. The disagreement on ii) is a consequence of the disagreement on i): MacIntyre thinks that there can be premises which are both descriptive and evaluative, and hence thinks that there can be valid deductive inferences from such descriptive premises to evaluative conclusions. I will discuss the issues surrounding this further in Chapter 2, Section 1.

⁷³ MacIntyre leaves open here whether he thinks there can be “pure” (entirely non-evaluative) descriptive claims, and whether these could or could not deductively imply evaluative conclusions. I will address this matter in Chapters 2 and 3.

Rather than intervening in this issue immediately, I will firstly focus on the issue of Johnston's position and ethical neutrality. I return to the questions on the relationship of descriptive and evaluative language at the end of this Section and in Chapters 2 and 3.

Having established that whether we will accept that we can deductively infer evaluative conclusions from descriptive premises depends on whether we think that any descriptive premises include one or more functional concepts, MacIntyre proceeds to the issue of whether Humean claims such as Johnston's⁷⁴ are ethically neutral. His answer is that they are not. For, as MacIntyre points out, the "Aristotelian tradition" – which I have been calling the "Classical teleological" tradition – maintains that at least one concept, "human being", is functional. *CT-1* and *CT-2* are examples of how Classical teleologists deploy this (in their view functional) concept in their moral reasoning.

The implication of this, as MacIntyre realises, is that Johnston's Fact/Value Dichotomy and Classical teleological views (or at least many of them) are incompatible. As he puts it, philosophers like Johnston have "taken for granted" that no moral arguments involve functional concepts, since otherwise, as we have seen, they would have realised that their claims were false, or perhaps, at best, half-truths about moral reasoning (since, if there are some moral arguments with concepts that are functional in the way MacIntyre describes, there can be valid deductive inferences from descriptive premises to evaluative conclusions).

But this means that Johnston takes for granted something that at least most kinds of Classical teleologist are committed to denying, namely that no moral arguments involve functional concepts. This results in the incompatibility between Johnston's supposedly neutral meta-ethical claims and at least some Classical teleological first-order views, i.e. those which maintain there are some moral arguments which include functional concepts.

⁷⁴ i.e. the view that we cannot have deductive inferences from descriptive premises to evaluative conclusions.

By my criterion, a position is not ethically neutral if it is incompatible with one or more first-order ethical views (see Ch. 1, Section 4, above). Hence, by my criterion, Johnston's claim that there can be no deductive inferences from descriptive premises to evaluative conclusions is not ethically neutral. I conclude on this basis that Johnston's purportedly neutral meta-ethical position is not ethically neutral at all, since it rules out most kinds of Classical teleological first-order view.

Conclusion to Section 4

In this Section, I have argued that Johnston's purportedly ethically neutral meta-ethical position is not, in fact, neutral, since it is incompatible with at least some versions of Classical teleological normative ethics. First of all, this reinforces the argument of Section 3, above, that Johnston cannot correctly claim it to be an advantage of his Emotivist meta-ethics that it is ethically neutral, while other meta-ethical positions, such as realist ones, do not. Secondly, by exposing which normative ethical views Johnston's position is incompatible with, we have made clear a vulnerability in Johnston's position. This vulnerability is that if the Classical teleological normative ethical views which conflict with Johnston's position can be well-defended, we will, in virtue of that, have reason to reject Johnston's position.

If we were to try to maintain Johnston's Fact-Value Dichotomy, then, we would need to reckon with its normative ethical commitments. Rather than being defensible purely on a neutral, linguistic basis, as Johnston suggests, we would now need to defend the normative ethical implications of the Fact/Value Dichotomy, if we were to maintain it. I therefore believe I have shown that defending the Fact/Value Dichotomy comes with an unexpected extra cost: one has to defend the denial of Classical teleological ethics which the Fact/Value Dichotomy necessitates. Whether this theoretical cost causes any serious problems for Johnston, however, I have not yet established. The reason is that if the relevant Classical teleological normative ethical views are not plausible, then Johnston

would be able to reply that his position, although not ethically neutral, only conflicts with normative ethical views which are, after all, not plausible.

What we found in this Section is that Johnston and the Classical teleologist come to blows on the issue of descriptive language, evaluative language, and the relationship between them. Specifically, as we saw above, Classical teleologists are committed to the idea that evaluative conclusions can be validly inferred from descriptive premises, and this turned out to be a result of their view that claims (such as “human beings are rational creatures”) can be both descriptive and evaluative. Johnston, by contrast, denies this possibility, arguing that claims cannot be simultaneously descriptive and evaluative. This brings out that, as we saw, Johnston must insist on a “splitting strategy” whereby he attempts to separate out the descriptive and evaluative elements of a sentence like “human beings are rational creatures” into logically independent claims. If Johnston can defend such a strategy successfully, this would give him a plausible defence against criticisms from Classical teleologists.

At this stage, I want to clarify that I have not yet attempted to offer any decisive counter-arguments against Johnston’s overall position. I have argued, in Section 3, that Johnston cannot make good on his claim to ethical neutrality, since no meta-ethical position can be ethically neutral. Moreover, I have argued that his meta-ethical position is inconsistent with (the main versions of) Classical teleological normative ethics. While the loss of the claim to neutrality removes one of the main purported advantages of Johnston’s position, it is not fatal to his main thesis that descriptive and evaluative language are of fundamentally different kinds (see Section 2, part ii), above), and it is this thesis which drives his Emotivism (see Section 1, above, for discussion of realism and Emotivism).

My arguments in this Chapter, then, while doing some modest damage to Johnston’s position by taking away his claim to neutrality, are more intended to clarify the battle-lines between Johnston and those, such as Anscombe and MacIntyre, who attack Johnston’s position from a Classical teleological direction. As we saw in this Section, Johnston and MacIntyre are committed to different analyses of descriptive and evaluative language. For

MacIntyre, concepts such as “watch” are both descriptive and evaluative, and hence to know that a good watch tells the time accurately is part of what it is to understand the very *concept* of “watch”. Johnston, by contrast, would deny this, arguing that a sentence like “your watch is grossly inaccurate in timekeeping” can be, and needs to be, separated out into separate descriptive and evaluative claims. This is how Johnston will respond to MacIntyre’s attempt to use the case of “watch” to give a counterexample to the view that we cannot deductively infer evaluative conclusions from descriptive premises.

So, as we saw in the case of *CT-1* and *CT-2* (the arguments for the conclusion that human beings ought to develop the intellectual virtues), Johnston will try to “split” certain premises into purely descriptive and purely evaluative claims, which are conceptually independent from one another. The question then arises, however, of whether these sentences are truly conceptually independent, as Johnston’s view requires them to be. In Chapter 2, I turn to these issues by investigating the phenomenon of apparently “thick” language in relation to Johnston’s views and the Fact/Value Dichotomy.

Chapter 2, Section 1: The Fact/Value Dichotomy

As I mentioned in Chapter 1, Section 2, I use “Fact/Value Dichotomy” as an umbrella term which covers more than one version of the basic idea that descriptive and evaluative language are different in kind. In this Section, I am going to give a clear statement of the stronger and weaker forms of the Dichotomy that I will be focusing on in this thesis. I will then discuss these different forms of the Dichotomy in relation to the recent literature, particularly the literature on “thick” language, which is often thought to pose a problem for the Fact/Value Dichotomy. I show why defenders of the Dichotomy have not been severely troubled by the apparent existence of “thick” language to date, expanding on some themes from the end of the last Chapter concerning the possibility of “separating” the descriptive and evaluative content of certain claims. I then use my analysis of this literature to identify a new strategy against the Fact/Value Dichotomy, which I will pursue and develop in Chapter 3.

Part i: The Humean Claim and the Dichotomous Claim: A Clarification

The stronger version of the Fact/Value Dichotomy, which is defended by Johnston (1989),⁷⁵ I call the Humean Claim.⁷⁶ This view is that *a descriptive claim can never be, nor imply, an evaluative claim*. In other words, according to this view, if a claim is descriptive, it *cannot* be evaluative, and if a claim is evaluative, it *cannot* be descriptive. It is a strong version of the Fact/Value Dichotomy, because it allows no overlap whatsoever between descriptive and evaluative claims: all claims that are either descriptive or evaluative can be one or the other, but not both. The Humean Claim also says that sets of premises which include only descriptive claims cannot entail conclusions with evaluative claims as their content.⁷⁷

⁷⁵ See Ch. 1, Sections 1 and 2.

⁷⁶ Whether this view is true to Hume is a controversial matter which I will not discuss in this thesis (see e.g. MacIntyre 1959).

⁷⁷ See Ch. 1, Section 2.

I want to make clear that, as I am understanding it, the Humean Claim is a view about *meaning*. This comes out very clearly from Johnston's (1989) discussion. Johnston's view is that the meaning of a descriptive claim like "this is yellow" is a "fundamentally different" kind of meaning to that of an evaluative claim like "this is good".

As we found in Chapter 1, Section 2, part iv), the Humean Claim is very hard to reconcile with moral realism. Realists agree that moral claims are in the business of reporting facts, and that there at least some true moral claims (Sayre-McCord 2005). But if there are moral truths which can be stated in language, it seems there must be true descriptive sentences which express those truths, which would have to be descriptive moral claims. For example, moral realists would surely agree that "killing an innocent human being for no reason is morally wrong" is a true descriptive claim.⁷⁸ However, they would also consider this claim to be evaluative, since it evaluates a particular kind of action as morally wrong. While realists of different kinds may disagree with each other about what *makes* this claim true, and what *kind* of description it is, they will, I think, agree at least that it is both a descriptive and an evaluative claim, and therefore reject the Humean Claim.

The incompatibility between the Humean Claim and realism is important here for two reasons. First, since this thesis critically engages with Emotivism, and thereby attempts to bolster the case for realism (see Chapter 1, Section 1, above), it is important for realists to give the best possible argument against the Humean Claim. In this Section, I will critically explore the Humean Claim, thereby setting up my argument against it in Chapter 3, Section 2. Secondly, however, the tension between the Humean Claim and realism, as we shall see in this Section, has influenced some realists to adopt a weaker (less radical) version of the Fact/Value Dichotomy. This raises the question of whether we ought to either follow these realists in accepting this weaker form of the Dichotomy or reject the Dichotomy in both forms. I will be making the case for the latter response in Chapter 3, but I first want to set out and discuss this more moderate version of the Fact/Value Dichotomy.

⁷⁸ An expressivist (e.g. Blackburn 1986, Gibbard 1990) would deny that it is true and argue instead that the speaker of such a claim expresses a sentiment, while an error theorist (e.g. Mackie 1977) would agree that the claim is descriptive but deny that it is true due to their view that there are no moral truths.

I call the weaker version of the Fact/Value Dichotomy the Dichotomous Claim. The Dichotomous Claim is similar to, but weaker than, the Humean Claim. According to the Humean Claim, any descriptive claims must be *purely* descriptive, in the sense that they must be entirely non-evaluative, and *vice versa*: any evaluative claims must be *purely* evaluative and entirely non-descriptive. This is why the Humean Claim is a strict version of the Fact/Value Dichotomy. The Dichotomous Claim waives this requirement. According to this Claim, it is *possible* for there to be “pure” (entirely non-evaluative) descriptive claims and “pure” (entirely non-descriptive) evaluative claims, but *not* all descriptive and evaluative claims are “pure” in this sense.⁷⁹ The Dichotomous Claim allows the possibility of some claims which are *both* descriptive and evaluative (such as, perhaps, the realist’s example of “killing an innocent human being is morally wrong”).

As I am defining them, both the Humean and Dichotomous Claims are about the *meaning* of descriptive and evaluative language. According to the Dichotomous Claim, there is a particular kind of meaning which is characteristic of descriptive claims, and a particular kind of meaning that is characteristic of evaluative claims, and these are of different kinds. The reason it is possible, on this view, for there to be “pure” descriptive claims is precisely that the meaning of such claims is of a distinct kind to the meaning of “pure” evaluative claims, and *vice versa*. However, this view allows that some claims – such as “thick” claims – combine the two kinds of meaning in a single sentence. As we will see later in this Section, on this view “thick” words are hence understood as possessing two distinct kinds of meaning at the same time (and a debate, we shall see, then ensues about whether these two kinds of meaning are “separable” from one another). The reason I nevertheless consider the Dichotomous Claim to fall under the umbrella of the Fact/Value Dichotomy is that it continues to consider the meaning of “pure” descriptive claims to be different in kind from the meaning of “pure” evaluative claims.

So, the advocate of the Dichotomous Claim will tell us, while Johnston’s view is too extreme in certain respects (i.e. in its denial that there is “thick” language), his basic

⁷⁹ This version of the Fact/Value Dichotomy is briefly mentioned, although not labelled, in Dancy (2013), cf. Harrison (2012). The distinction is also anticipated by Lovibond (1983) (whose view I critically discuss in Chapter 2, Section 3, below).

intuition that there is a sort of semantic difference between descriptive and evaluative language is still correct. What I want to establish in this thesis is whether we ought to give this “Dichotomous” response to Johnston, or whether there is some more fundamental mistake in Johnston’s account than the advocate of the Dichotomous Claim realises.⁸⁰

While I will eventually argue against the Dichotomous Claim, the view does have several interesting advantages, which means it requires more careful critical scrutiny than it generally receives in the literature.⁸¹

The advocate of the Dichotomous Claim, as I define it, can remain neutral on whether there *are* any claims which are both descriptive and evaluative at the semantic level (I will make what I mean by this more specific momentarily). But allowing this possibility, logical space opens up between this Claim and the Humean Claim, which does not allow this possibility. Relating back to the dispute we considered between MacIntyre and Johnston about concepts such as “watch” and “human being” in Chapter 1, Section 4, this means that the Dichotomous Claim is compatible with MacIntyre’s view that those concepts are both descriptive and evaluative, while Johnston’s view, as we saw, is not. So whether we accept the Humean Claim or the Dichotomous Claim (or neither) makes a significant difference to our meta-ethics.

One apparent advantage of the Dichotomous Claim, which gives more reason to investigate it further, is that it is much easier to reconcile with moral realism than the Humean Claim. Indeed, one of the primary defenders of the Dichotomous Claim, Parfit (2011), combines the Dichotomous Claim with moral realism. Since the Dichotomous Claim allows the possibility of claims which are both descriptive and evaluative, it can accommodate the moral realist’s view that claims like “killing an innocent human being is morally wrong” are both descriptive and evaluative. Many realists, such as Parfit, avoid the apparent problems with the Humean Claim by simply embracing the attractive-seeming Dichotomous Claim instead. But, as I shall argue in Chapter 3, Section 2, we should reject the Dichotomous Claim too. This will lead us to a form of Naturalist realism different to

⁸⁰ Cf. my discussion in Chapter 1, Section 1, above.

⁸¹ The remainder of Ch. 2 will critically examine how the Dichotomous Claim has been treated in recent literature.

those which Parfit (2011) considers.

In fact, it is worth considering Parfit's commitment to the Dichotomous Claim in more detail here, since I orient my discussion around his position later, in my argument against the Dichotomous Claim. Parfit insists that "the deepest theoretical questions are about the purely normative concepts and claims that we can express with [thin evaluative words]" (2011: 266). Thus, meta-ethicists such as Parfit attach theoretical primacy to purportedly "thin" evaluative words. This has a significant effect on Parfit's meta-ethical argumentation, as I explore in Chapter 3, Section 3. Since Parfit centres his analysis on purportedly "thin" evaluative words, this leads him not to consider a possible version of meta-ethical Naturalism which *denies* that thin descriptive language is possible (see Chapter 3, Section 3, below). I will argue, however, that such a form of Naturalism is highly plausible, and satisfactorily answers one of Parfit's own main objections to Naturalism.

It is not so much that Parfit (2011) or anybody else gives an argument for the Dichotomous Claim. It is more that the Dichotomous Claim looks to be supported by a *prima facie* examination of descriptive and evaluative words. As much as it looks like some words are thick, like "brave" and "lewd", Parfit thinks there are also clear examples of "thin" words, i.e. "purely" descriptive and evaluative words. His purported examples are words like "thin", "square" and "tall" on the descriptive side, and words like "right", "ought" and "duty" on the evaluative side (2011: 266). Much like in the Humean Argument presented by Johnston (see Chapter 1, Section 2, above), Parfit has the impression of a fundamental distinction between words that are used to describe the things we perceive with our senses, and words used to evaluate things as positive or negative.

However, Parfit's less radical version of the Fact/Value Dichotomy seems to be much more plausible and easier to defend than the Humean Claim, at least at first sight. Parfit's way of supporting it, by noting the apparent difference between words like "thin", "square" and "tall" on one hand, and "right", "ought", and "duty" on the other, seems particularly convincing. Since he does not endorse the Humean Claim, he does not need to argue that

all our descriptive and evaluative language is “thin”, but only that *some* of it is “thin”. And this argument appears to be adequately supported by his examples of words like “thin” and “ought”, since it does seem quite convincing, at first glance, that these words are of different kinds to one another.

For this reason, I believe that, if we are to oppose the Dichotomous Claim, the main challenge is to answer Parfit’s contention that the Dichotomous Claim is supported by a *prima facie* examination of descriptive and evaluative words. This is a challenge I take up in Chapter 3, Section 2. Although I accept that the Dichotomous Claim is *prima facie* plausible, I will argue that it rests on a mistaken view about meaning. To my knowledge, no other argument which pursues this approach against Parfit has yet been produced in the existing literature, so this will be one of the primary original contributions of this thesis.

The purpose of this Chapter, however, is to pave the way for that later argument, by discussing these two forms of the Fact/Value Dichotomy in relation to recent literature. First, in the remainder of this Section, I discuss my distinction between the Humean Claim and the Dichotomous Claim with reference to the literature on thick concepts. Doing so will help clarify further the apparent strengths and weaknesses of each version of the Dichotomy, and help us identify the right way to resist the “splitting strategy” that I suggested Johnston could deploy against MacIntyre’s counterexamples in Chapter 1, Section 4. In the next two Sections of the Chapter, I consider two objections to the Fact/Value Dichotomy from the existing literature (those of Putnam 2002 and Lovibond 1983) and explain why I find them inadequate.

Before proceeding, we should note that if the Dichotomous Claim falls, the Humean Claim will fall with it. The reason for this is that the Humean insists on the possibility of “pure” descriptive and evaluative claims. This is why he insists that, if a claim is either descriptive or evaluative, it must be one or the other and cannot be both. If that turns out *not* to be a possibility, both the Humean and Dichotomous Claims would be false, since they both insist that this is a possibility. The Humean Claim and Dichotomous Claim both entail that “pure” descriptive and evaluative claims are possible; the difference is that the latter allows

that “thick” claims, which are both descriptive and evaluative, are also possible. I will now consider why this feature of the Dichotomous Claim arguably makes it more plausible than its Humean cousin.

Part ii: Thick Concepts and the Humean and Dichotomous Claims

Since the Dichotomous Claim entails less than the Humean Claim, it is comparatively easier to defend. One important issue, which demonstrates this, is the way both these versions of the Fact/Value Dichotomy deal with the purported phenomenon of “thick” language. I will now explore this issue, with a view to setting up my arguments in Chapter 3.

One reason many philosophers find the Humean Claim implausible is that they believe that some language is “thick”, in the sense that it is both descriptive and evaluative. They think that “thick” language can be used to construct “thick” claims, which are also both descriptive and evaluative, yielding counterexamples to the Humean Claim. In this respect, such philosophers would likely side with MacIntyre and against Johnston on the specific issue of the nature of concepts like “watch”: those philosophers who believe in “thick” concepts would allow that, at least in principle, “watch” could be both descriptive and evaluative (*pace* Johnston). Other examples of seemingly “thick” words which combine description and evaluation include “brave”, “cruel”, “lewd”, and “rude” (see e.g. Dancy 1995, Dancy 2013, Vayrynen 2013, Kirchin 2013).

So, for example, the claim “Mulan is brave” seems to be thick because it *describes* Mulan’s character (it gives us concrete information that, for instance, she will not cow from danger easily), but also *evaluates* her character (“brave” is a term of praise and so the claim in question here evaluates Mulan’s character as, at least in some respect, good). Moreover, as we shall see in more detail later in this Section, many such philosophers analyse such “thick” words in terms of two aspects of their meaning, or two different kinds

of “content”, one corresponding to the “descriptive” part of the word’s meaning, and one corresponding to the “evaluative” part of the word’s meaning (see e.g. Elstein and Hurka 2009, for discussion see Vayrynen 2017: Section 3.1, Section 4, although for a different view⁸² see Vayrynen 2013).

If this is right, then “Mulan is brave” and other such “thick” claims are both descriptive and evaluative, and would therefore violate the semantic thesis of the Humean Claim, according to which a single claim cannot combine the kind of meaning possessed by descriptive claims and the kind of meaning possessed by evaluative claims (e.g. Johnston 1989, see Chapter 1, Section 2, part ii), above).

I also want to introduce the terminology of “thin” language, which will help us better understand and discuss the Humean and Dichotomous Claims. Many philosophers (e.g. Parfit 2011, Vayrynen 2013, 2017) use the term “thin” to denote words and claims which, unlike “thick” claims, which are both descriptive and evaluative, are “purely” (non-evaluatively) descriptive or “purely” (non-descriptively) evaluative. So, to recap using this new term, the Humean is committed to saying that *all* descriptive and evaluative claims are “thin” at the semantic level: each claim that is either descriptive or evaluative can only have *one* of the kinds of meaning (either the meaning characteristic of “descriptive” claims or the meaning characteristic of “evaluative” claims) and not both. The Dichotomous Claim, by contrast, allows that some claims can be semantically “thick”, but also maintains that it is possible to have semantically “thin” descriptive and evaluative claims.

One important problem for Johnston (1989) is that he says nothing about this apparent phenomenon of “thick” language and how it seems to threaten his core thesis. At least initially, one might think that this is because his Humean position cannot satisfactorily account for, or explain, the existence of apparently “thick” language. While, as we will see, matters are not quite so straightforward, one significant attraction of the Dichotomous Claim over the Humean Claim is that it seems to offer a much better way of handling the

⁸² Vayrynen (2013), as I shall address below, rejects the idea that “thick” concepts have two different kinds of content, on the grounds that the “evaluative” nature of these concepts is to be located not at the semantic level, but at the level of pragmatics.

apparent phenomenon of thick language.

We should note, however, that the existence of apparently-thick language is not, on its own, a decisive consideration against the Humean Claim. The reason for this is that not all philosophers agree that words like “brave” and “cruel” are really “thick” in the sense that they are, in terms of their meaning, both descriptive and evaluative. For example, Vayrynen (2013) argues that thick words are, in terms of their meaning, no different to “plain” descriptive words like “table”, “chair”, “red”, etc. He acknowledges that words such as “brave” are often used to evaluate, but he argues that this is a matter of the *pragmatics* of the word, rather than the semantics. So, on Vayrynen’s view, at a semantic (logical) level, there are no thick concepts. (I will not be assessing Vayrynen’s view here, since my criticism of the Fact/Value Dichotomy does not depend on or appeal to the view Vayrynen attacks, i.e. the bipartite analysis of thick terms as possessing two distinct kinds of meaning (see part iii of this Section and Chapter 3, below).

Since the Humean Claim states that, on a semantic level, claims which are either descriptive or evaluative must be one or the other and not both, Humeans can allow that we *use* so-called thick words to both describe and evaluate, but must (like Vayrynen) deny that this dual descriptive-and-evaluative aspect of thick words is present at the semantic level. It is at the semantic level that a Humean maintains that all language is “thin”. I understand this to mean that, for the Humean, descriptive words and claims have one kind of meaning, and evaluative words and claims have another kind of meaning (see Johnston’s (1989) view as explained in Chapter 1, Section 2). This is not necessarily to say that there are such things as “descriptive meaning” and “evaluative meaning”, which could confuse matters,⁸³ but rather to say that there are at least two different kinds of meaning, and that descriptive words and claims have one kind while evaluative words and claims have another.

⁸³ These terms are sometimes used in the literature (see e.g. Kovesi 1967), but I believe use of these terms adds an unnecessary complication. As I explained, I do not see why anybody in the debate (including advocates of either form of the Fact/Value Dichotomy) need to commit to there being “descriptive meaning” or “evaluative meaning”. The two kinds of meaning that they attribute to descriptive and evaluative claims could be broader than applying to descriptive and evaluative claims only, for instance.

So, to recap, it is not so easy to refute the Humean Claim simply by presenting examples of apparently “thick” words. The Humean could respond to the phenomenon of apparently-thick words by accounting for their apparent dual descriptive-evaluative nature by appeal to some non-semantic feature or features they possess, much like Vayrynen (2013). Hence, one route towards resisting the Humean Claim would be to challenge whether Humeans can give a good account which meets this criterion. This is not the route I pursue in this thesis, however, and for that reason I do not need to go into whether Vayrynen’s (2013) pragmatics-based account of the “thickness” of certain concepts is plausible. Instead, I will resist the Humean Claim by attacking an assumption which is common to both the Humean and Dichotomous claims, namely that there can be semantically “pure” descriptive language is possible (see Chapter 3, below).

I also want to note that the fact that the Humean Claim cannot be easily dismissed by appeal to thick concepts is one reason why the arguments I developed against Johnston in Chapter 1, Sections 3 and 4, are important. They give us reason to doubt Johnston’s claims, and seek an alternative to his position, independently of what the correct view is about thick and thin language, since we found that his position cannot deliver on ethical neutrality, and that it conflicts with some normative ethical views. However, we do not yet have a knock-down argument against Johnston’s Humean Claim.

Although the apparent trouble caused for the Humean by the apparent existence of thick concepts could potentially be adequately dealt with, then, it nevertheless seems like a comparative advantage of the Dichotomous Claim, over the Humean Claim, that it does not *require* us to deny that there are semantically thick concepts. Adopting the Dichotomous Claim instead of the Humean Claim would hence relieve us of the burden of arguing, like Vayrynen (2013), that there are no semantically thick concepts.

If we opt for the Dichotomous Claim instead, the apparent existence of thick concepts will not be a problem for us, because the Dichotomous Claim is perfectly compatible with the existence of semantically thick concepts. Equally, if Vayrynen (2013) is right and there are no thick concepts, that is compatible with the Dichotomous Claim, too. The Dichotomous

Claim, let us recall, is that it is *possible* to have “pure” (entirely non-evaluative) descriptive claims and “pure” (entirely non-descriptive) evaluative claims, but it also allows that there may be language which is both descriptive and evaluative (i.e. “thick” claims). It is worth considering precisely how the Dichotomous Claim handles the issue of thick concepts, to help us check whether adopting the Dichotomous Claim would genuinely help us be able to take a more easily defensible position than the Humean Claim on this issue.

Advocates of the Dichotomous Claim can disagree on the issue of how widespread “thick” words and claims are in ordinary language. Putnam (2002: 40) suggests that, for Williams (1985), all, or nearly all, our present ordinary language is in fact “thick”. However, at the same time, Williams (as interpreted by Putnam, at least) also believes that it is possible, in principle, to construct semantically “thin” language, even if no ordinary language at the present time is “thin”. By contrast, Parfit (2011) thinks that much of our present ordinary language is “thin”. (I will discuss Parfit’s reasons for believing this in more detail later in this Section). In a similar way, an advocate of the Dichotomous Claim could say that all or most of our ordinary language is semantically “thin” but concede that semantically “thick” claims are possible. The Dichotomous Claim is flexible here, because it merely asserts that both semantically thin and semantically thick claims are *possible*.

So those such as Parfit (2011), who accept the Dichotomous Claim, allow that there may be “thick” concepts, which are both descriptive and evaluative. Since, as I stated above, the Dichotomous Claim is a *semantic* claim about the meaning of descriptive and evaluative claims, the Dichotomous Claim must understand thick words as somehow combining the semantic qualities which it takes to be characteristic of descriptive words with the semantic qualities it takes to be characteristic of evaluative words. The fact that the Dichotomous Claim is a *semantic* claim about the possibility of semantically “thick” language is important since it differentiates its stance on thick concepts from the Humean Claim, for instance. It is open to the Humean to say that certain concepts are *non-semantically* “thick”, for instance by adopting Vayrynen’s (2013) view, mentioned above, but not to say that there are or could be *semantically* thick concepts. But what exactly is it to say that there are *semantically* thick concepts? In what sense do these words combine two different kinds of meaning, and how does the meaning of semantically “thick” concepts differ from

the meaning of non-thick concepts? We need answers to these questions if we are to make clear what is involved with accepting the Dichotomous Claim. To provide those answers, I will now survey the recent debates in the literature on thick concepts, especially the dispute between “Separabilist” and “Inseparabilist” accounts of thick concepts (see e.g. Vayrynen 2013, 2017 and the papers in Kirchin 2013).

Part iii: Separabilism about Thick Language

Recall how in Chapter 1, Section 4, we considered that Johnston might try to answer MacIntyre’s view that we can validly infer evaluative conclusions from exclusively descriptive premises by insisting that we can and should “split” certain sentences (which appear both descriptive and evaluative) into separate logically independent claims which are “purely” descriptive and “purely” evaluative. This “Separabilist” approach to thick concepts is, as we are about to see, a way of developing this proposal as a way of understanding apparently “thick” language quite generally. Hence, if Separabilism about thick concepts is viable, this would help Johnston give a strong response to the challenge from MacIntyre considered in Chapter 1, Section 4. I will come back to address MacIntyre’s challenge to Johnston in the last part of this Section.

According to the “Separabilist” view (for examples see Burton 1992, Tappolet 2004, Elstein and Hurka 2009), thick claims like “Mulan is brave” are composed of two distinct semantic elements which can, at least in principle, be analysed down into two self-standing claims, one of which would be non-evaluatively descriptive⁸⁴ (see Vayrynen 2017), and one of which would be evaluative. In conducting this process of separation, we would start with a claim which possessed both the kind of meaning which characterises descriptive claims and the kind of meaning which characterises evaluative claims, and end up with two claims which each possess only one of these kinds of meaning, with each possessing a

⁸⁴ It is significant that Vayrynen uses the phrase “non-evaluatively descriptive”, especially because from the meta-ethical realist’s point of view, there is no reason why description may not itself be evaluative. To be clear, then, it is open to an advocate of the Fact/Value Dichotomy to draw the line between evaluation and *non-evaluative* description (cf. Parfit’s view as discussed in Chapter 3, Section 3, below).

different kind. Or, in other words, the Separabilist view is that “thick” claims have both “descriptive” content and “evaluative” content, and that it is possible to *reduce* the “thick” claim into two distinct claims, one of which only has the descriptive content, and one of which only has the evaluative content (see Vayrynen 2017: Section 3.1).⁸⁵

To give a concrete example, a Separabilist may argue that “Mulan is brave” can, in principle, be broken down into two distinct claims:

- 1) Mulan is Brave-D.
- 2) Mulan is Brave-E.

Here, “brave-D” is supposed to stand for the descriptive “content” of “brave”, while “brave-E” stands for the evaluative “content” of brave. So, perhaps, “Mulan is Brave-D” could stand for the “non-evaluative descriptive” features assigned by the word “brave”, such as that Mulan does not typically shy from danger easily, while “brave-E” could stand for the evaluative features assigned by “brave”, such as “virtuous”, “of good character”, etc.

The advocate of the Humean Claim may well be keen to take up the Separabilist view, for the following reason. The Humean could argue that the Separabilist analysis proves that “Mulan is brave” and other apparently “thick” claims are *not*, despite appearances, single claims, but rather placeholders for two distinct claims, one of which is descriptive and one of which is evaluative. If correct, this would show that sentences like “Mulan is brave” are not counterexamples to the Humean Claim that no claims can be simultaneously descriptive and evaluative (semantically speaking). For in this case, the Separabilist analysis would have shown that sentences like “Mulan is brave” are not single claims, but

⁸⁵ If reader thinks this is not a satisfactorily clear statement of the Separabilist thesis, I am inclined to agree: I will later be arguing that we cannot make sense of this idea that the descriptive and evaluative semantic “elements”, or “content”, of thick concepts, are different in kind from one another. I hence regard the possible need for further clarification of the Separabilist thesis to be a problem for Separabilists (and other advocates of the Fact/Value Dichotomy) and not one for me.

rather placeholders for two different claims, which fit within the Humean strictures. The Humean would thereby show that sentences like “Mulan is brave” are not counterexamples to the Humean Claim after all. Such a strategy is not without precedent: structurally, this proposal is similar to Russell’s (1905) view that “The King of France is bald” stands for three distinct claims: that there is a King of France, that exactly one entity is the King of France, and that entity has the property of baldness. (Note the similarity here between Johnston’s attempt to split a single sentence into separate “pure” descriptive/evaluative claims in Chapter 1, Section 4). So Separabilism seems like it might help Humeans give a plausible response to the apparent problem for their view created by apparently “thick” language.

Vayrynen (2017) also claims, correctly, that it need not trouble the Separabilist that there may be no word in natural languages corresponding to e.g. the purely descriptive semantic component or concept of “brave”, for it seems possible for the concept to be there without its happening to have its own word.⁸⁶ Compare how Aristotle notes certain virtues and character traits happen to lack a single word that refers to them in *Nicomachean Ethics*, Ch. 2 (Aristotle 2000 [340 BC]). Plainly, then, it is possible for there to be a concept, and for us to contingently lack a word to express it, which is good news for the Separabilist analysis of thick language.

Even if we are not convinced that sentences like “Mulan is brave” and other thick claims really stand for two distinct claims, however, we may still be attracted to Separabilism. For to be a Separabilist about thick words, we need only claim that it is possible *in principle* to separate thick words and claims into their constituent “evaluative” and “non-evaluative descriptive” components. This point hinders part of Putnam’s argument against the Fact/Value Dichotomy, because Putnam seems to think that the failure of Separabilists to reduce certain thick words (e.g. “cruel”) into their descriptive and evaluative “components” speaks against Separabilism (Putnam 2002: 38). But if this failure is only a *contingent* failure, this point will not do much damage to Separabilism or the Humean Claim. As long as it is plausible that a “Separation” into “pure” descriptive and evaluative

⁸⁶ See Vayrynen, *ibid.*, Section 3.1.

elements is possible, the Humean Claim will remain plausible too.

Part of the reason this Separabilist view is important in the context of this thesis, then, is that it seems to help the Humean Claim account for the apparent phenomenon of thick language. Such language, the Humean can say, *appears* to be semantically “thick”, but is consistent with the idea that a descriptive claim can never be, nor imply, an evaluative claim, because upon analysis such “thick” sentences are shown to stand for independent, “pure” descriptive and evaluative claims. But Separabilism is also significant here because it can be adopted by the advocate of the Dichotomous Claim too.

If we accept the Dichotomous Claim, we do not *have* to argue that “Mulan is brave” is merely a placeholder for two different claims, since we allow that there can be semantically “thick” claims, which are both descriptive and evaluative. However, if we accepted the Dichotomous Claim, we could still accept Separabilism, by maintaining that it is possible to separate “thick” claims such as “Mulan is brave” into a “pure” (non-evaluative) descriptive claim and a “pure” (non-descriptive) evaluative claim. This would cohere with the idea behind the Dichotomous Claim that it is possible to have claims which are “thinly” descriptive and “thinly” evaluative, and that such “thin” claims have meanings which are different in kind from one another. According to this combination of the Dichotomous Claim and Separabilism, then, “thick” claims combine two distinct sorts of meaning, and it is, at least in principle, possible to “separate” these two sorts of meaning by separating out the single “thick” claim into two “thin” claims. On this view, thick claims are understood to have a bipartite semantic structure: they are hybrids which combine, in a single claim, both the kind of meaning possessed by “thin” descriptive claims and the different kind of meaning possessed by “pure” evaluative claims (Vayrynen 2017). (The bipartite nature of this view, incidentally, is my reason for using the name “Dichotomous Claim”).

While the Dichotomous Claim is compatible with Separabilism, there is no pressure for advocates of the Dichotomous Claim to embrace Separabilism. If it turned out to be impossible in principle to “separate” thick claims into “thin” components, the advocate of

the Dichotomous Claim could simply respond by saying that while thick claims may be impossible to “separate” into thin claims, this is no reason to say that there *are* no thin claims. They can still equally well maintain that there are thin descriptive and evaluative claims, as we saw Parfit (2011) does (see part i) of this Section, above), only adding that thick claims cannot be separated into two or more thin ones.

Once again, the Dichotomous Claim seems to enjoy an attractive flexibility which the Humean Claim lacks, owing to its ability to either embrace, reject or remain neutral on Separabilism, as the evidence requires. One way of proceeding, then, would hence be to press the Humean view on whether Separabilism or Vayrynen’s approach (outlined earlier in this Section) are ultimately defensible responses to the problem for the Humean Claim caused by the apparent existence of thick language. If we concluded that they are not, then perhaps the most natural stopping point would be to accept the Dichotomous Claim, which is not committed to either Separabilism or Vayrynen’s approach. And indeed, as we shall see in part v) of this Section, accepting the Dichotomous Claim will lead us to say Johnston is wrong to rule out that evaluative conclusions can be validly deductively inferred from descriptive premises (see Chapter 1, Sections 2 and 4, above).

However, rather than making the case for the Dichotomous Claim over the Humean Claim, I instead want to critically examine a crucial common feature of the Humean and Dichotomous Claims, namely that they both consider it possible for there to be “thin” descriptive claims, which have one kind of meaning, and some “thin” evaluative claims which have a different kind of meaning. This feature, which both forms of the Fact/Value Dichotomy share, turns out to be mistaken, I shall argue in Chapter 3.

To sum up, then, both Humeans and advocates of the Dichotomous Claim can embrace the “Separabilist” view about thick words and claims. This brings out that the mere existence of apparently-thick sentences like “Mulan is brave” does not, on its own, pose a fatal problem for either the Humean or the Dichotomous Claim. The Humean can argue that such sentences are not truly “thick” in the semantic sense, perhaps taking up the Separabilist account, on which sentences like “Mulan is brave” are placeholders for two

different “thin” claims, or instead adopting Vayrynen’s approach of explaining the descriptive and evaluative nature of such sentences without appeal to their semantics. Meanwhile, an advocate of the Dichotomous Claim can accept that claims like “Mulan is brave” are thick in the semantic sense, but argue that this does not threaten the basic Johnstonian thesis that the meaning of thin descriptive claims is different in kind from the meaning of thin evaluative claims. On this view, Johnston (1989) is wrong to rule out the existence of semantically “thick” claims, but still nevertheless right that there are, fundamentally, two different kinds of meaning in play here, one of which is characteristic of “thin” descriptive claims, and the other of which is characteristic of “thin” evaluative claims.

My subsequent argument, in Chapter 3, Section 2, will be that even this Dichotomous position concedes too much to Johnston, supporting my earlier contention that many meta-ethicists concede too much to Johnston and that this damages their understanding of moral language (cf. my description of the argument of this thesis in Chapter 1, Section 1, above). For now, though, let us recap what we have learned about the Humean Claim, the Dichotomous Claim, and Separabilism. I will then consider the place of “Inseparabilism” about thick concepts in the debate and use that discussion to explain some of the motivations behind the argument of Chapter 3.

As we have seen, an advocate of the Dichotomous Claim can understand “thick” claims as hybrids which happen to possess both “kinds of meaning”; an idea which they can, perhaps, cash out in terms of “thick” claims having both descriptive “content” and evaluative “content”. This answers the questions I raised at the end of part ii) of this Section about what it means, specifically, to say that there are, or can be, *semantically thick* claims.

Moreover, Separabilism seems to offer advocates of the Humean Claim, like Johnston, a good way of responding to the controversy we considered in Chapter 1, Section 4, about arguments from descriptive premises to evaluative conclusions. As we saw in that Section, Johnston will try to respond to the Anscombe/MacIntyre-style counter-examples by

“splitting” certain premises into component “thin” claims, thereby preserving his view that evaluative conclusions cannot be entailed by descriptive premises only. Since Separabilism, then, appears to have favourable implications for the Fact/Value Dichotomy, it is worth considering the primary alternative to Separabilism, namely Inseparabilism, and considering the implications of that view for the Humean and Dichotomous Claims.

Part iv: Inseparabilism about Thick Language

What may appear to be threatening, at least *prima facie*, to the Fact/Value Dichotomy is the “Inseparabilist” view about thick words and claims. I will now outline the Inseparabilist view, before explaining why most versions of Inseparabilism do not, after all, pose a problem for the Fact/Value Dichotomy. However, this will lead me to identify and motivate an alternative route of attack against the Humean and Dichotomous Claims, which involves questioning the possibility of semantically “thin” descriptive language (for more detail see Chapter 3, Section 1 and for my argument, see Chapter 3, Section 2).

Vayrynen (2017: Section 3.2) describes Inseparabilism as the view that the meanings of thick words/concepts “involve both evaluation and non-evaluative description without this being a matter of combining constituent evaluative and non-evaluative contents” (for examples of this view see Dancy (1995, 2013) and Harcourt and Thomas (2013)). The Inseparabilist challenge to the Johnstonian idea that there are two different kinds of semantic “content” present in thick language appears to threaten the Fact/Value Dichotomy, which, as we saw in the last part of this Section, relies upon that idea to explain the existence of apparently “thick” language. However, the challenge for Inseparabilism, Vayrynen (2017: Section 3.2) says, is “to pin down more precisely how thick terms are both evaluative and descriptive in ways that are continuous with, respectively, thin evaluative concepts and non-evaluative concepts”, while still resisting the Separabilist claim that these are distinguishable elements which can be pulled apart.

The Inseparabilist strategy, then, as understood by Vayrynen, is to accept the distinction between “thick” and “thin” language, and, on the basis of this model of “thin” descriptive and evaluative language, to show how “thick” words are both descriptive and evaluative. But for this view to remain Inseparabilist, it must also deny that thick language has two distinguishable semantic elements which can be separated and form self-standing “thin” claims, since otherwise it would, by definition, just collapse into Separabilism.

For short, I am going to call the kind of Inseparabilism Vayrynen describes “conventional Inseparabilism” (Later in this part, I will distinguish this from another kind which I call “radical Inseparabilism”).

I will now express some worries about conventional Inseparabilism. Although none of my later arguments in Chapter 3 depend on these worries being valid, this will help motivate my choice of strategy against the Fact/Value Dichotomy.

It seems to me that the task Vayrynen sets for Inseparabilists is curiously quixotic. Vayrynen asks Inseparabilists to accept the notion of “thin” descriptive and evaluative language, i.e. language which only has one of the “kinds of meaning” posited by the advocate of the Fact/Value Dichotomy, then asks Inseparabilists to show thick terms are descriptive and evaluative “in ways that are continuous with... thin evaluative concepts and non-evaluative concepts”. This, I suppose, can only mean that thick language must be understood as possessing *both* the kind of meaning characteristic of “thin” descriptive language and the kind of meaning possessed by “thin” evaluative language. If we accept this understanding of thick language, however, how will we maintain that the descriptiveness and evaluativeness of thick language is *not* a matter of its “combining constituent evaluative and non-evaluative contents”, as Vayrynen thinks Inseparabilists must? To my ears, this seems like demanding that Inseparabilists show that thick concepts have two distinct semantic components without allowing that thick concepts have two distinct semantic components, and clearly this demand is impossible to satisfy.

Moreover, if Inseparabilists accept that there can be “thin” descriptive and evaluative claims, each of which possess a fundamentally different kind of meaning, this would raise the question: why, then, would a Separabilist analysis of thick language *not* be possible? If “thin” descriptive and evaluative claims are capable of independent substantiation, does this not increase the plausibility of it being at least *possible in principle* to “separate” a thick claim into its “thin” semantic components? For this reason, I believe that accepting the possibility of “thin” language, as conventional Inseparabilism does, significantly improves the prospects for Separabilism. This is perhaps why Inseparabilists are often left appealing to our present inability to conduct a Separabilist analysis of certain thick words and claims to support their position (see e.g. Putnam 2002: 38 and cf. McDowell 2000; I discuss Putnam in more detail in Chapter 2, Section 2.) But this is a futile argument, because it does not threaten the Separabilist view that a Separabilist analysis is *possible*. In other words, our present inability to conduct such an analysis in some cases may be contingent. And as long as this possibility can be maintained, the Fact/Value Dichotomy will continue to be safe: the mere contingent fact that we cannot always separate sentences into their “thin” components does not genuinely threaten the Fact/Value Dichotomy.

Although none of my arguments henceforth depend on this, I am, for the reasons just mentioned, sceptical about the prospects for conventional Inseparabilism. This should lead us to consider what alternative options there are for those sympathetic to Inseparabilism (such as McDowell 2000, Lovibond 1983, Putnam 2002, Chappell 2013, Dancy 1995, 2013, Harcourt and Thomas 2013), and who suspect that the phenomenon of apparently-thick language causes more trouble for the Fact/Value Dichotomy than generally thought.⁸⁷ (At the same time, if conventional Inseparabilism *can* be rendered successful, this is perfectly consistent with my own positive views in Chapter 3. I am not aiming to either defend or refute conventional Inseparabilism in this thesis).

My alternative proposal is that Inseparabilists should challenge the notion of “thin” language, rather than uncritically accept it, as Parfit (2011) and Vayrynen (2013, 2017)

⁸⁷ The philosophers mentioned here are not univocal in their views on thick concepts or the Fact/Value Dichotomy. Sections 2 and 3 of this Chapter examine this literature, with particular focus on Putnam and Lovibond, who offer the most direct arguments against the Fact/Value Dichotomy.

suggest we should. In this respect I am building on Dancy's (1995: 268) denial that Inseparabilists should think "there are two "really" distinct elements [of thick concepts] which by a pseudo-chemical reaction somehow become indistinguishable from each other". That is, the Inseparabilist should not uncritically accept the idea that so-called "thin" descriptive language possesses one distinct type of meaning and that so-called "thin" evaluative languages possesses another, different-in-kind type of meaning and that these two distinguishable semantic elements are somehow combined together in "thick" language. Rather, we should instead question whether there are two genuinely distinct kinds of meaning in play here at all. If there are not, then the very idea of semantically "thin" descriptive or evaluative language is confused. Hence, Inseparabilists need not aim "to pin down more precisely how thick terms are both evaluative and descriptive in ways that are continuous with, respectively, thin evaluative concepts and non-evaluative concepts", as Vayrynen suggests. We should not aim to explicate thick terms in such a way that their descriptive and evaluative natures are shown to be "continuous" with thin concepts, because the notion of "thin" descriptive concepts and language is, I shall argue in Chapter 3 spurious, and skews the debate in favour of Separabilism. I call this alternative approach "radical Inseparabilism".

If the notion of "thin" language is indeed spurious, then Separabilism is not tenable. Since Separabilism insists it is possible to "separate" thick language into thin components, it depends on the idea that there can be "thin" descriptive language. The arguments of this thesis, in other words, should be good news for Inseparabilists, as long as Inseparabilists are willing to drop what I called "conventional Inseparabilism" and take up my more radical position, which challenges the core of the Fact/Value Dichotomy.

If we are to find fault with the Dichotomous Claim, then I believe it is much more promising to pursue this "radical Inseparabilist" option than to pursue "conventional Inseparabilism". As we saw above, conventional Inseparabilism is set up in a way that is hospitable to the Dichotomous Claim. Vayrynen's way of defining the task for Inseparabilism, for instance, takes the notion of thin language for granted and asks Inseparabilists to explain how thick language is descriptive and evaluative in a way that is "continuous with thin descriptive and evaluative language".

Indeed, conventional Inseparabilism is entirely compatible with the Dichotomous Claim, for reasons mentioned in part iii), above: if it is merely that thick language cannot be separated into thin language, but that thin language is still either actual or at least possible, then the Dichotomous Claim is still true. If, by contrast, thick language cannot be separated into thin language *because* the very notion of thin (descriptive, or evaluative, or both) language is spurious, then the Dichotomous Claim should be rejected, since its claim that “thin” language is possible would then be problematic.

This shows why making the case for Inseparabilism is not sufficient to reject the Fact/Value Dichotomy. It may help to undermine the Humean Claim, but it certainly does not help undermine the Dichotomous Claim – at least, unless one adopts radical Inseparabilism. (This is one reason why distinguishing clearly between the Humean and Dichotomous Claims, as I have in this Section, is so vital for understanding these issues clearly).

Hence, questioning the notion of “thin” language is a promising route towards unravelling the Fact/Value Dichotomy. As I will argue in Chapter 3, Section 2, we have good reasons to judge the notion of “thin” descriptive language, which underpins the Fact/Value Dichotomy, to be confused. I support this position by developing wrongly neglected arguments from Kovesi (1967) and Meynell (1971).

Let me sum up what I have done so far in this Section. I firstly distinguished between a stronger and weaker version of the Fact/Value Dichotomy, one which I call “the Humean Claim” (which is defended by Johnston (1989)), and one which I call “the Dichotomous Claim” (which is defended by Parfit (2011)). At first sight, it appears the Humean Claim is less plausible than the Dichotomous Claim, because the apparent existence of semantically thick claims seems to provide counterexamples to the Humean Claim but not the Dichotomous Claim. However, I then considered how Humeans could respond to this problem, with the aim of showing that the apparent existence of semantically thick concepts and claims is not sufficient to refute the Humean Claim. I have canvassed ways

for both Humeans and advocates of the Dichotomous Claim to account for the phenomenon of apparently-thick claims like “Mulan is brave”, focusing on Separabilist and Inseparabilist strategies for understanding such claims. Rather than directly tackling the Humean or Dichotomous ways of understanding thick claims, I proposed rather to scrutinise a commitment that we found the Humean and the advocate of the Dichotomous Claim to share: the idea that there can be semantically “thin” descriptive and evaluative language. If this brick can be dislodged from the Fact/Value Dichotomy’s pyramid, the entire structure will crumble.

In the remainder of this Chapter, I will consider some objections to the Fact/Value Dichotomy offered by other recent philosophers, namely MacIntyre (1981), Lovibond (1983), Putnam (2002). Using my distinction between the Humean Claim and Dichotomous Claim, as well as the considerations above about how advocates of these views can account for supposedly “thick” claims, I will explain why these previous objections to the Fact/Value Dichotomy are unsuccessful. However, as we shall see, each philosopher’s approach does include certain insights that shape my alternative, improved objection to the Fact/Value Dichotomy.

Part v: The Humean Claim, the Dichotomous Claim and Anscombe/MacIntyre Counterexamples

Now we have distinguished between the Humean Claim and the Dichotomous Claim, I want to discuss these claims in relation to the supposed counterexamples to the Fact/Value Dichotomy offered by Anscombe (1958) and MacIntyre (1981), which I mentioned in Chapter 1, Section 4.

In that Section, the point of discussing those counterexamples was to draw out some normative ethical commitments implicit in Johnston’s position. But there is a separate issue concerning these alleged counterexamples, which relates more closely to the reason

Anscombe and MacIntyre presented them in the first place: do these counterexamples not simply show that the Fact/Value Dichotomy is mistaken? Anscombe and MacIntyre both seem to think so. Is there, then, any point in conducting a fuller argument against the Fact/Value Dichotomy in the remainder of this thesis? Or do Anscombe and MacIntyre's counterexamples settle the issue already?

Since the Humean Claim says that descriptive premises alone can never entail evaluative conclusions, MacIntyre's counterexample appears to threaten the Humean Claim:

P1) This watch is grossly inaccurate and irregular in time-keeping.

C) This watch is a bad watch (see MacIntyre 1981: 58 and the discussion in Chapter 1, Section 4, above).

Since this is a perfectly good inference, and P1) appears to be descriptive while C) appears to be evaluative, it seems like a counterexample to the Humean Claim. (We should note, however, that the example does not threaten the Dichotomous Claim, and this shows some of the value of explicitly distinguishing these different forms of the Fact/Value Dichotomy: again, the Dichotomous Claim looks more robust against critical attack than the Humean Claim). As I observed in Chapter 1, Section 4, the problem with such Anscombe/MacIntyre-style counterexamples to the Humean Claim is that, in order to make such arguments logically valid, we would need to add a suppressed premise:

P1) This watch is grossly inaccurate and irregular in time-keeping.

P2) If a watch is grossly inaccurate and irregular in time-keeping, it is a bad watch.

C) This watch is a bad watch.

This unpacked version of the argument, however, does not threaten the Humean Claim, since Humeans, like Johnston, allow that a descriptive premise can combine with an evaluative premise to yield an evaluative conclusion, which is what our unpacked argument appears to do. Since Anscombe and MacIntyre do not explicitly consider this, at this point it is uncertain whether their purported counterexamples against the Humean Claim succeed.

As I remarked in Chapter 1, Section 4, the crucial question is whether the supposedly “suppressed” premise (in this case P2, “If a watch is grossly inaccurate and irregular in timekeeping, it is a bad watch”) is just unpacking the meaning of one or more of the concepts in the first premise. If it is, then it would be a mistake to think of these “premises” as logically independent, as the Humean needs to construe them (if they are going to ward off the threat that their view will be disproved by counterexample). If P2 is just unpacking the meaning of P1, then P2 is not a true “suppressed premise”, since it would already be implicit in the first premise. (To understand this point, the reader should refer back to the “James is a bachelor” example in Chapter 1, Section 4). In this case, we would need to think of P1 as both descriptive *and* evaluative, and semantically thick, since it would have turned out to have “evaluative content” (content which is represented by P2). But this would be inconsistent with the Humean Claim.

The reason that it is not easy to turn this into a winning objection against the Humean Claim is that whether “evaluative” content like P2 can be part of the meaning of “descriptive” claims like P1 is precisely what Humeans like Johnston (1989) call into question (see Chapter 1, Sections 1 and 2, above). So Anscombe and MacIntyre cannot just assume that P2 merely unpacks the meaning of P1 without begging the question against Johnston. For that reason, their purported counterexamples do not, on their own, refute the Humean Claim, as I said in Chapter 1, Section 4. For MacIntyre’s counterexample-based objection to the Humean to be convincing, then, we would need some independent reason to think that “suppressed premises” such as P2 really are just explanations of the meaning of their corresponding premises like P1, much in the same way as “a bachelor is an unmarried man” (the “suppressed premise” in our bachelor-based example argument in

Chapter 1, Section 4) explains the meaning of “James is a bachelor”. In Chapter 3, I will give some reasons for thinking of examples like P2 are exactly that.

As I have just argued, the purported counterexamples from Anscombe (1958) and MacIntyre (1981) do not, on their own, defeat the Humean Claim. They are even less effective against the Dichotomous Claim. Since advocates of the Dichotomous Claim allow that there can be thick language, they can respond to the Anscombe/MacIntyre examples by simply allowing that certain concepts featuring in the examples, such as “watch”, are thick. So for example, an advocate of the Dichotomous Claim could say that “watch” is a thick concept, and therefore the concept “watch” has certain evaluative content, such as its being part of the concept of “watch” that a good watch tells time accurately. By allowing this, the advocate of the Dichotomous Claim could accept that, in MacIntyre’s “watch” example, P2 really is just unpacking the meaning of P1, and could hence accept that MacIntyre’s example is really one of a descriptive premise (P1) entailing an evaluative claim. This would mean giving up the Humean Claim.

Note however that on this response P1 would be, semantically, *both* descriptive and evaluative: if “watch” is thick, then a claim about a watch or watches will be thick too (i.e. both descriptive and evaluative). In the case where we have thick descriptive premises in an argument, then, the advocate of the Dichotomous Claim is free to say that these premises can entail evaluative conclusions. This is significantly different from the Humean response, since the Humean does not accept the possibility of semantically thick language. The advocate of the Dichotomous Claim will hence be keen to stress that Anscombe/MacIntyre counterexamples do not threaten her position.

A question remains, however, about whether the advocate of the Dichotomous Claim would allow that one or more “thin” descriptive premises could entail (thin or thick) evaluative conclusions. It seems more likely that they would not: for since, on their view, thin descriptive claims have no evaluative content, how could they possibly entail an evaluative conclusion? We would now be back with the original Humean thought that one cannot “magic up” evaluative content from nowhere (see Chapter 1, Sections 1-2, above).

If there is no evaluative content in the premises, the thought goes, there cannot be any in the conclusion – if the argument is to be logically valid, anyway. The reason that some Anscombe/MacIntyre counterexamples are possible is that some descriptive concepts also have evaluative content, but if we are dealing with “thin” description, there will be (according to the Dichotomous Claim) no evaluative content and hence no possibility of entailing an evaluative conclusion.

Thus, the advocate of the Dichotomous Claim will, I think, respond to Anscombe and MacIntyre by saying that their examples showcase the phenomenon of thick language, but do not threaten the idea that semantically thin language is possible either. Hence, taking everything I have argued in this part together, we can say that, on their own, the Anscombe/MacIntyre examples do not make either the Humean Claim or the Dichotomous Claim implausible. If we can find no problems with these Humean/Dichotomous responses to the examples, we will have to conclude that the examples just do not threaten the Fact/Value Dichotomy. Yet perhaps we should not rest content with the Dichotomous Claim’s response to these cases just yet. In Chapter 3, I will argue that the Dichotomous Claim’s notion of semantically thin language has not yet been adequately explained. Since the idea that “thin” descriptive premises cannot entail evaluative conclusions is clearly dependent on whether there can⁸⁸ be semantically thin language, critically scrutinising whether the notion of semantically thin language has been adequately explained, as I do in Chapter 3, shows a way of making Anscombe/MacIntyre counterexamples threatening to the Fact/Value Dichotomy. I will therefore return to Anscombe/MacIntyre-style counterexamples, and what they portend for the Fact/Value Dichotomy, at the end of Chapter 3.

⁸⁸ I am using the modal term “can” here deliberately. Even if all present ordinary language is semantically thick, the advocate of the Dichotomous Claim could still say that, if there *were* any semantically thin descriptive premises (in the future, for example), then these could not entail evaluative conclusions (see this Section, above). Hence, we need to tackle the modal nature of the Dichotomous Claim if we are really to make progress against it. Merely arguing that present ordinary language is all semantically thick (like Chappell 2013, for instance) will not be sufficient to escape the Fact/Value Dichotomy.

Chapter 2, Section 2: Putnam's Rejection of the Fact/Value Dichotomy

In this Section, I am going to set out Putnam's reasons for rejecting the Fact/Value Dichotomy (or, more precisely, taking himself to reject the Fact/Value Dichotomy). I will then critically respond to Putnam's arguments, explaining why I do not believe his arguments are successful. Along the way, I will try to draw some lessons to bear in mind for my own response to the Fact/Value Dichotomy. Specifically, this will relate to recognising the need to target both the Humean and Dichotomous Claims if we are to avoid the Dichotomy.

Part i: A Summary of Putnam's Argument

As we saw in the previous Section, the Fact/Value Dichotomy, in both its stronger and weaker forms, relies on the idea that there are two distinct kinds of meaning, one of which is possessed by "thin" descriptive claims, and one of which is possessed by "thin" evaluative claims. To adapt a remark from McDowell (2000: 38), the advocate of the Fact/Value Dichotomy thus needs to provide accounts of these two kinds of meanings. They must explain why these two kinds of meaning are genuinely different from one another (cf. Chapter 3, Sections 1 and 2, below). On top of this, they must also make plausible that one of these kinds of meaning really is possessed by "thin" descriptive claims and the other is possessed by "thin" evaluative claims, unless they want to argue that "thin" language is only a possibility rather than a present actuality. In the latter case, they would need to say something about what the two sorts of possible meaning they have in mind are, in order to substantiate their view that it is possible for there to be "thin" descriptive language (with one kind of meaning) and "thin" evaluative language (with a different kind of meaning).

Should the advocate of the Fact/Value Dichotomy fail in this regard, he will have failed to make his claims intelligible. It is futile to insist that "thin" descriptive claims and "thin"

evaluative claims possess different kinds of meaning to one another if one cannot say anything about these two kinds of meaning, and why they are different in kind from one another (cf. McDowell 2000: 38 and, for a related but slightly different point, McManus 2013: 19). This points us towards a potential argumentative strategy against the Fact/Value Dichotomy, one which is pursued, in different ways, by Putnam (2002), Lovibond (1983) and the present author (in Chapter 3, Section 2). If we can expose the accounts of meaning used to prop up the Fact/Value Dichotomy as erroneous, the Dichotomy's thesis that there are two kinds of meaning in play here will fall.

Johnston's (1989) account is one attempt to make good on this need to give an account of the two different kinds of meaning, to make the Fact/Value Dichotomy viable. This is one reason I judged it important to evaluate Johnston's account, as I did in Chapter 1, above, but unfortunately Putnam (2002) does not discuss Johnston's views in his book.

Another preliminary remark concerning Putnam's (2002) view is that he does not explicitly distinguish, as I have, between the stronger form of the Fact/Value Dichotomy, the Humean Claim, and the weaker form, the Dichotomous Claim (see Chapter 2, Section 1, above). As a result, it is not always clear whether Putnam is trying to target both of these views, or only one of them, as we will see later in this Section. As I will explain shortly, this makes Putnam's arguments less useful than they might otherwise have been for opposing the Fact/Value Dichotomy.

Putnam (2002) presses the advocate of the Fact/Value Dichotomy on the issue of the supposed two kinds of meaning which, according to advocates of the Dichotomy, are characteristic of descriptive and evaluative language, respectively. Putnam asks: what led advocates of the Fact/Value Dichotomy to think that descriptive claims and evaluative claims have two different kinds of meaning in the first place? He poses this question because he thinks that the Fact/Value Dichotomy rests upon a mistaken account of the meaning of descriptive, factual claims (like "the cat is on the mat" and "all crows are black").

Specifically, Putnam argues that the Fact/Value Dichotomy is at home alongside a “pictorial semantics” (2002: 15), and he identifies Hume as the founding father of such views:

Hume’s criterion for “matters of fact” presupposed what might be called a “pictorial semantics”. Concepts, in Hume’s theory of the mind, are a kind of “idea”, and “ideas” are themselves pictorial: the only way they can represent any “matter of fact” is by *resembling* it (not necessarily visually... ideas can also be tactile, olfactory [etc.]) (Putnam 2002: 15).

According to this “pictorial semantics”, meaningful descriptive⁸⁹ statements in some way represent, or stand for, certain states of affairs. So, on Hume’s theory, as understood by Putnam, the ideas in the proposition “there is a red chair in the library” in some manner resemble, and thus represent, the state of affairs that is a red chair being in the library. If the relevant state of affairs obtains in reality, then the proposition corresponding to it is true, otherwise it is false. The meaning of the descriptive statements is thus understood in terms of the picture(s), or state(s) of affairs, which is it supposed to represent.⁹⁰ Putnam’s suggestion is that this “pictorial semantics” account of descriptive claims gives rise to the Fact/Value Dichotomy, because it is hard to see how evaluative claims could possibly be “descriptive” in this way.

It is worth noting that the basic “pictorial semantics” approach can be cashed out in various different ways, many of which Putnam considers (see his 2002: Chs. 1-2). So for instance, Putnam considers both Hume’s version of the view and the views of logical positivists such as Carnap (2002: 22-23). It is best, therefore, to understand “pictorial semantics” as a family of approaches with similar core elements. But the core elements of such views which are most relevant for our present purposes were those I recapped in the previous paragraph.

⁸⁹ Or, on certain views, simply meaningful statements *tout court*. Some philosophers including Ayer (1936) have argued that only descriptive empirical claims are meaningful, and that claims of other sorts (such as, arguably, moral claims) are hence nonsensical.

Putnam's discussion of "pictorial semantics" in relation to the Fact/Value Dichotomy brings us back to the *Humean Argument* I considered in Chapter 1, Section 2. The Humean Argument has two simple steps. First, it says that if some claim is to be descriptive, it must conform to the "pictorial semantics" model we have just sketched (Putnam 2002: 15). Secondly, it denies that evaluative claims can conform to this model. As Putnam puts it, on this view, for an evaluative claim to be descriptive, "the property of virtue [and other moral properties] would have to be *picturable* in the way that the property of being an apple is picturable" (2002: 15, emphasis in original). As we found in Chapter 1, Section 2, part iii), advocates of the Fact/Value Dichotomy, such as Johnston, do not find it plausible that "virtue" or other moral properties could be picturable in this way (Johnston 1989: 142).

In fact, Putnam is inclined to agree with Johnston on that specific point: *if* we accept the pictorial semantics account of descriptive claims, he says, then the Fact/Value Dichotomy follows naturally (2002: 15). Since it is so hard to imagine, for instance, what state of affairs a claim like "eating meat is wrong" could represent, by contrast to claims like "the cat is on the mat", which are easily picturable, we get the strong impression that there is a fundamental difference in kind between descriptive and evaluative claims. Hence, Putnam thinks, the Fact/Value Dichotomy is at home alongside the "pictorial semantics" account of descriptive claims.

However, Putnam thinks that we ought to reject the "pictorial semantics" on which the Fact/Value Dichotomy and the Humean Argument, on his view, rest. There are several reasons why Putnam rejects the "pictorial semantics" account of descriptive language, with the common theme being his idea that "so much of our descriptive language is a living counterexample" to it (2002: 26).

For example, Putnam thinks that it does not adequately account for psychological descriptive claims which include one or more "folk psychological" concepts, such as "she intended to kill him" (2002: 25-26). He argues that the "pictorial" model does not work for such sentences, yet they are plainly descriptive claims. He considers that the most plausible candidate for a "picture" that such a psychological descriptive claim could "represent"

would be a certain “brain-state” in the person being described (2002: 26). However, Putnam thinks it implausible that such mundane psychological claims are about brain-states, writing that

the idea that whenever I describe someone as cruel, or as irritated, or as delighted, I am committed to a “theory” according to which there is a “brain-state,” either physical or computational in character, such that all cruel (or all irritated, or all delighted) people are in that brain-state, and no one who is not cruel (or irritated, or delighted) is in that brain-state is not scientific at all, but mere science fiction. (2002: 26).

So folk-psychological descriptive sentences, according to Putnam, are not adequately captured by pictorial semantics. In other parts of his discussion, Putnam also suggests that the pictorial semantics model may have trouble with (arguably-) descriptive claims which involve ascriptions of necessity (2002: 33) and which feature abstract theoretical terms such as “charge” (2002: 29).

While I am sympathetic to Putnam’s arguments against pictorial semantics, I am not going to scrutinise them in detail here. The reason is that I am happy to grant Putnam that the pictorial semantics cannot account properly for the meaning of the range of descriptive claims in our language – to accept, in his words, that there are “living counterexamples” to pictorial semantics as an account of our descriptive language. I shall be arguing that, even if we grant Putnam this, his argument is effective only against the Humean Claim, and not against the Dichotomous Claim. For this reason, Putnam does not help us escape the Fact/Value Dichotomy as I am defining it. I argue against both the Humean Claim and the Dichotomous Claim in Chapter 3.

Part ii: Problems with Putnam’s Argument

The trouble with Putnam’s argument is that an advocate of the Fact/Value Dichotomy could accept his claims *without* giving up the Fact/Value Dichotomy. The advocate of the

Fact/Value Dichotomy could respond to Putnam in at least two different ways. First, he could challenge Putnam's argument that the account does not adequately capture the meaning of different kinds of descriptive claims. But, as I stated above, I am willing to grant Putnam a free pass against this objection. The second response, which is the one I shall discuss and develop here, is to accept Putnam's claim that pictorial semantics does not capture the meaning of all descriptive claims, but to say that this is, after all, compatible with the Fact/Value Dichotomy (specifically the Dichotomous Claim).

The idea behind this second response is to say that while the pictorial semantic model does not apply to all descriptive claims, it does nevertheless apply to a restricted class of descriptive claims. Recall the distinction I made in the previous Section, between the stronger version of the Fact/Value Dichotomy (the Humean Claim) and the weaker version (the Dichotomous Claim). According to the Dichotomous Claim, it is *possible* to have claims which are "purely" or "thinly" descriptive, i.e. claims which have an entirely different kind of meaning to that possessed by evaluative claims. So if *some* descriptive claims in actual language *do* fit the pictorial semantics model, the advocate of the Dichotomous Claim has what he was looking for: at least *some* examples of "thin" descriptive language. For the Dichotomous Claim, it makes no matter if some descriptive claims do not fit into the pictorial mould, as Putnam claims. Putnam's view is hence entirely consistent with the Dichotomous Claim. Putnam's core claim, we saw, is that there are "living counterexamples" to the pictorial semantics account. But the existence of counterexamples to that account does not establish that *no* descriptive language is adequately captured by it.

Indeed, from examining Putnam's text, it seems clear that he intends only to target the more radical form of the Fact/Value Dichotomy (i.e. the Humean Claim) with his arguments. For instance, he writes "that... thick ethical concepts are counterexamples to the idea that there exists an absolute fact/value *dichotomy* has long been pointed out" (2002: 35, emphasis in original). This wording suggests that when Putnam talks about the "Fact/Value Dichotomy" he is not including the Dichotomous Claim, as I am in this thesis, because the existence of thick concepts is plainly *not* a counterexample to the Dichotomous Claim, which allows the existence of thick concepts and language (see

Chapter 2, Section 1, above). Hence, if Putnam included the Dichotomous Claim as a version of the “Fact/Value Dichotomy”, he would not say that thick concepts are a counterexample to that Dichotomy.

Putnam then considers responses that have been made to the idea that thick concepts are counterexamples to the Fact/Value Dichotomy, and explains why he does not consider any of these responses satisfactory (2002: 34-40). Hence, for Putnam, the existence of thick concepts provides good reason to reject what he calls “the Fact/Value Dichotomy”. But as we saw in the previous Section, the existence of thick concepts is no threat to the Dichotomous Claim, even if it is a sort of threat to the Humean Claim. Since, when Putnam criticises what he calls the “Fact/Value Dichotomy”, his arguments cause problems for the Humean Claim but not the Dichotomous Claim, I believe it is most plausible to say that by “Fact/Value Dichotomy” Putnam is referring to the Humean Claim, but not the Dichotomous Claim.

To take another example which supports this interpretation, one of Putnam’s most important arguments against the “Dichotomy” is that thick concepts cannot be “separated” into “pure” descriptive and evaluative components (2002: 38). As we saw in Section 1 of this Chapter, though, the Fact/Value Dichotomy, as I am defining it, does not depend on the possibility of so separating thick concepts: even if thick concepts and/or claims cannot be separated into “thin” elements, this would not rule out that some language is thin. This means that Putnam’s argument against Separabilism cannot touch the Dichotomous Claim, even if it speaks against the Humean Claim.

Since Putnam does not explicitly distinguish between the Humean Claim and Dichotomous Claim, as I have, it is tricky to work out which version(s) of the Dichotomy Putnam does and does not intend to criticise. Outlining his view, Putnam writes:

If we *disinflate* the fact/value dichotomy what we get is this: there is a distinction to be drawn (one that is useful in some contexts) between ethical judgments and other sorts of judgments. This is undoubtedly the case, just as it is undoubtedly the case that there is a distinction to be drawn (and

one that is useful in some contexts) between *chemical* judgments and judgments that do not belong to the field of chemistry. *But nothing metaphysical follows from the existence of a fact/ value distinction in this (modest) sense.* (2002: 19, emphasis in original).

While this clarifies that Putnam wants to avoid metaphysical claims and that he is happy to accept some sort of distinction between ethical and non-ethical judgments, this alone does not differentiate his view from Parfit's (2011) view, for instance. Clearly Putnam accepts some "disinflated", non-metaphysical version of the Fact/Value Dichotomy, but what exactly this amounts to is left unclear. For instance, does Putnam accept that some language is, or at least could be, semantically thin? He is silent on this question, which makes it hard to know where Putnam stands on the Dichotomous Claim.

Due to this, I do not find Putnam's arguments suitable for resisting the Fact/Value Dichotomy (as I define it). The Dichotomous Claim requires different treatment to the Humean Claim, for reasons I discussed in the previous Section of this Chapter, and Putnam does not directly supply the necessary treatment, perhaps because he does not explicitly distinguish the Humean and Dichotomous Claims.

It is, however, reasonable to ask whether Putnam's arguments could be adapted to make them effective against the Dichotomous Claim. Perhaps the failure of the pictorial semantics model when applied to folk-psychological claims, for instance, which Putnam emphasises, casts doubt on whether it is a plausible account of meaning of *any* descriptive claims at all, even the ones it appears best suited to explain (like "the cat is on the mat"). However, it is not clear whether or why this should be so. It seems that we cannot rule out that different kinds of descriptive claims may have distinct kinds of meaning. So it is fairly plausible, *prima facie* at least, that something like the pictorial semantics picture might apply to a restricted class of descriptive claims, but not others. If this were the case, then pictorial semantics might be the correct account of seemingly-"thin" descriptive claims, even if it fails to capture other kinds descriptive language. This would reinforce the Dichotomous Claim, by supporting the idea that "thin" descriptive language has a special kind of meaning which is different in kind to that of evaluative claims, even if not all descriptive claims have that kind of meaning.

Unless Putnam is able and willing to give an argument for why we should expect there to be a single universal account of the meaning of all descriptive claims (and he does not indicate that he is), he cannot rule out this possibility, and hence cannot rule out the Dichotomous Claim.

Putnam is, I think, correct to identify that certain questionable assumptions about meaning, and especially the meaning of descriptive claims, underpin the Fact/Value Dichotomy. In this respect, my argument in Chapter 3 is similar to Putnam's. Otherwise, however, the arguments Putnam puts on the page are not threatening to the Dichotomous Claim (and indeed may never have been intended to be so). To adapt Putnam's arguments to make them threaten the Dichotomous Claim, we would need to completely reject pictorial semantics as an account of descriptive claims quite generally, to rule out the possibility that "thin" descriptive claims are covered by the pictorial model. But Putnam's objection to pictorial semantics, based on ruling it out as an account of *certain kinds* of descriptive language (like folk-psychological descriptive claims), is ill-suited to this task. Finally, Putnam's arguments against Separabilism do not help against the Dichotomous Claim, for reasons mentioned in the first Section of this Chapter.

We could respond to this in two different ways. We could simply say that the fact Putnam's argument does not threaten the Dichotomous Claim is really a merit, not a drawback. After all, the Dichotomous Claim appears plausible, and, as Parfit (2011) says, well-supported by a *prima facie* inspection of descriptive and evaluative language (see Chapter 2, Section 1). I, however, will be taking a different direction. I do not think we should accept the Dichotomous Claim, for reasons I shall present in Chapter 3, Section 2. Therefore, I do not think Putnam's argument on its own is adequate, even if it is effective against the Humean Claim. With my later argument, I aim to undermine both the Humean and Dichotomous Claims. The fact that Putnam's argument targets only one of these forms of the Fact/Value Dichotomy shows why it is so essential to explicitly distinguish between these forms, as I did in Section 1 of the present Chapter. I now turn to an argument from the existing literature which *does* genuinely target the Dichotomous Claim.

Chapter 2, Section 3: Lovibond’s Rejection of the Fact/Value Dichotomy

In this Section, I will explain, and then critically respond to, Sabina Lovibond’s (1983) rejection of the Fact/Value Dichotomy. Unlike Putnam’s (2002) argument, Lovibond’s argument targets both the Humean Claim and the Dichotomous Claim. However, I will then expose several problems with Lovibond’s argument. Most important of these, I will argue, is that Lovibond’s view depends on very strong essentialist claims about language and meaning, which I judge very hard to defend. I therefore conclude that opponents of the Fact/Value Dichotomy would do well to find alternative arguments, which do not depend on such strong, controversial claims. This motivates and shapes my later argument in Chapter 3, Section 2.

Part i: A Summary of Lovibond’s Argument

First, though, I will present Lovibond’s view. Lovibond (1983: 36) is critical of the attempt to “[drive] a wedge between factual and evaluative meaning”, which sounds reminiscent of the views of the present author and Putnam (2002), for example. Notably, Lovibond does *not* use the labels “Humean Claim” and “Dichotomous Claim”, which I introduced in Section 1 of this Chapter, nor does she distinguish explicitly between those two views. One other difference in the language Lovibond uses, compared to the present author, is that she speaks of “factual” and “evaluative” meaning. I have avoided using these terms, because I do not think advocates of the Fact/Value Dichotomy need to commit to the idea that there are such things as “factual” (or “descriptive”)⁹¹ meaning and “evaluative” meaning, even though these terms are sometimes used in the literature (see e.g. Kovesi 1967: 3-4). Rather, the advocate of the Fact/Value Dichotomy commits to either the possibility of semantically “thin” descriptive and evaluative language, and the idea that “thin” descriptive language has a different kind of meaning to “thin” evaluative language. Whether these two kinds of

⁹¹ Lovibond would, I think, have done better to say “descriptive” than “factual” here, but in any case, she intends to attack the same Dichotomy that this thesis attacks, as I explain shortly.

meaning are properly called “descriptive” (or “factual”) and “evaluative” is a further matter, on which advocates of the Fact/Value Dichotomy, as I am defining it, may disagree. (I will not be passing judgment on that further matter in this thesis).

Despite these minor differences between Lovibond (1983) and the present author on terminological points, Lovibond’s aim to prevent a “wedge” being driven between what she calls “factual meaning” and “evaluative meaning” is clearly relevant to, and in the spirit of, my own aim, which is to reject the Fact/Value Dichotomy. For example, in the opening chapters of *Realism and Imagination in Ethics*⁹² (1983) Lovibond makes clear that she is opposing non-cognitivist meta-ethical theories and the ideas about language which underpin them, which is similar to MacIntyre’s (1981) attack on “Emotivism” (see Chapter 1, Section 1, above).

While Lovibond thus conceives the target of her arguments slightly differently to the present author, I believe Lovibond’s core strategy can be adapted and rephrased so that it targets the Fact/Value Dichotomy as I am understanding it. While the advocate of the Fact/Value Dichotomy may not need to “drive a wedge” between “factual” and “evaluative” meaning, they still need to “drive a wedge” between two kinds of meaning, one which they think is possessed by thin descriptive claims, and the other which they think is possessed by thin evaluative claims (cf. my comment on McDowell (2000) at the beginning of Chapter 2, Section 2). So, in what follows, I will be setting out Lovibond’s view as it can be used to try to deal with the task at hand, namely evaluating the Fact/Value Dichotomy as it is defined in in Chapter 2, Section 1.

Like Putnam (2002), Lovibond (1983) also attempts to give a genealogy of the Fact/Value Dichotomy. Lovibond’s genealogy is somewhat different from Putnam’s, since she introduces some distinctions which do not feature in Putnam’s discussion. For Lovibond, the “wedge” between the two supposedly different kinds of meaning arises from the

⁹² Lovibond’s subsequent book *Ethical Formation* (2002) does not, so far as I can tell, include new material relevant to the specific issue of the Fact/Value Dichotomy, which is why I do not discuss *Ethical Formation* in this thesis.

“empiricist” idea that some claims have a special property which other claims lack, namely “reference to an objective reality” (1983: 25). The idea is that, in virtue of this special property, the claims which possess it have a different kind of meaning to those which do not. In Lovibond’s view, the “empiricist” considers factual descriptive claims (e.g. “smoking causes cancer”) to have this special property of “reference to an objective reality”, but thinks that evaluative claims (like “eating meat is wrong”) do not, and this difference is supposed to generate the semantic difference between these two kinds of claim (cf. Johnston 1989 and my discussion of his view in Chapter 1, Section 2).

What, exactly, does Lovibond mean by “reference to an objective reality”? To explain the “empiricist” view, Lovibond says “It is essential to the empiricist idea of [descriptive] language... that as soon as our definitions have been stated, what we can correctly say is *out of our hands*” (1983: 18, emphasis in original). Furthermore, on the “empiricist” view, “given our initial definitions of natural-language expressions, the configuration of objects in the world will determine the truth-value of propositions *independently of any mediation by us*” (1983: 19, emphasis in original). It seems that Lovibond has something similar in mind to Putnam, when he discusses “pictorial semantics” (see Chapter 2, Section 2, above). The basic idea seems to be that factual descriptive claims refer to particular states of affairs, and so if the “configuration of objects in the world” is right, the claim comes out as true. This “empiricist” account is hence one on which the truth-conditions for a claim are intimately connected with the meaning of the claim (for critical discussion of this kind of view see Dummett 2006: Ch. 4). As with the empiricist pictorial semantics, the “empiricist” view Lovibond considers, then, holds that descriptive claims are somehow specially connected with particular states of affairs, and that this relation (however it is cashed out in detail) somehow shapes or constitutes the meaning of the claim in question.

Lovibond classifies this “empiricist” view as a “metaphysically heterogeneous” view of language (1983: 27-8, 36-7). By this, Lovibond means that, for the empiricist, there are claims – like factual descriptive claims – which have a privileged metaphysical relation to reality, due to their special property of “reference to an objective reality”, but also claims which *lack* this privileged metaphysical status. The typical “empiricist” non-cognitivist, Lovibond thinks (1983: 21), takes moral claims to belong to the latter category, since they

do not think that moral claims refer to objective reality (otherwise they would be realists of some sort) and they do not think the meaning of moral claims should be analysed in terms of their truth-conditions (since, as non-cognitivists, they deny that moral claims are truth-apt).

The “metaphysically heterogeneous” view Lovibond identifies is therefore relevant to the Fact/Value Dichotomy as defined in this thesis. For if this metaphysically heterogeneous view is right, Lovibond explains, it could be used to explain the supposed difference in the kind of meaning possessed by thin descriptive claims and the kind possessed by thin evaluative claims: the metaphysically heterogeneous view of language, she says,

gives rise to the idea that there are two contrasting kinds of meaning that words can have: on one hand “descriptive” or “cognitive” meaning; on the other, “evaluative” or “emotive” meaning” (1983: 21).

Lovibond is suggesting that the “empiricist” explains the meaning of thin descriptive claims by appeal to their supposed property of “reference to an objective reality”, which thin evaluative claims (according to the advocate of the Dichotomy) lack. Thus, if thin evaluative claims are meaningful at all, their meaning must be different in kind (cf. Johnston 1989 and my explanation of the “Humean Argument” in Chapter 1, Section 2).⁹³

If the “metaphysically heterogeneous” view of language divides claims into those with the special metaphysical status of “reference to an objective reality” and those lacking it, then a “metaphysically homogeneous” view, according to Lovibond, would be one which draws no distinction between claims which supposedly have such a special metaphysical status

⁹³ If the gap here between the claim that thin descriptive claims have a special metaphysical property and the claim that they have a special kind of meaning is a problem, this is a problem for the advocate of the Fact/Value Dichotomy who wants to establish their Dichotomy this way, not a problem for Lovibond or the present author. I will not pass judgment on whether that *is* a problem for the Dichotomy in this thesis, since I think the motivation behind the Dichotomous Claim comes from a different source (see Chapter 3, Section 2, below).

and those which do not. Lovibond thinks we should reject the metaphysically heterogeneous view of language in favour of the metaphysically homogeneous view:

“reference to an objective reality” cannot intelligibly be set up as a target which some propositions... may hit, while others fall short... if something has the grammatical form of a proposition, then it *is* a proposition” (1983: 26, emphasis in original).

When Lovibond says that “if something has the grammatical form of a proposition, then it *is* a proposition”, she indicates her disagreement with the non-cognitivist idea that sentences like “telling the truth is morally good” are not really propositions, since they are not truth-apt. Most non-cognitivists acknowledge that denying that such sentences are propositions is counterintuitive, since they *appear* to be propositions, just like “telling the truth is very common” is a proposition. So by stating that things which *appear* to be propositions *are* propositions, Lovibond signals her intent to avoid the typical non-cognitivist manoeuvre of saying that sentences like “telling the truth is morally good” appear to be propositions, but are not genuine propositions (cf. Smith 1994).

Indeed, one of the main advantages Lovibond takes her view to have is that she thinks it is better supported by ordinary language than the metaphysically heterogeneous view. For Lovibond, there is theoretical merit in allowing that “if something has the grammatical form of a proposition, then it *is* a proposition”. She finds it implausible that a distinction can be drawn between “parts of assertoric discourse which do, and those which do not, genuinely *describe* reality” (1983: 36). (In denying this distinction, she is not trying to undermine the distinction between true and false claims, but rather than distinction between descriptive and non-descriptive claims, see 1983: 26).

Rather, for Lovibond it seems more plausible, and more in-keeping with ordinary language, not to attempt to artificially drive a wedge between some parts of assertoric discourse and others, as non-cognitivists, for instance, try to do. In this respect, Lovibond’s defence of realism fits with the pattern I noted in Chapter 1, Section 1, which is that

realism seems, at least typically, to fit better with ordinary language than non-cognitivist, “Emotivist” views, like expressivism.

Diamond (1996: 221) worries that the fact Lovibond takes up a “metaphysically homogeneous” view of language means that she will be unable to account for the *diversity* of our language and the multiplicity of roles played by different words and sentences. Indeed, Diamond argues, ignoring or forgetting about this diversity is one of the main wellsprings of philosophical confusion (Diamond 1996: 221, Wittgenstein 1953: Section 122, cf. Fogelin 1996, Chapter IX). But, even granting Diamond that we need to account for the diversity of roles played by different kinds of words and claims in our language, her objection against Lovibond misses its target, because it rests on a misunderstanding of Lovibond’s view. Lovibond explicitly distinguishes between “metaphysical” and “phenomenal” claims about language:

The homogeneity of language is not, of course, asserted [by Wittgenstein] at the phenomenal level (the level at which we “describe language-games, cf. *PI* 486), for there are manifestly “countless different kinds of use of what we call “symbols”, “words”, “sentences”” (*PI* 23). It is asserted at the metaphysical level – the level at which empiricism drives a wedge between factual and evaluative meaning (1983: 36).

So, in claiming language to be “metaphysically homogeneous”, Lovibond is not, after all, denying that there are “countless different kinds of use of... “symbols”, “words”, “sentences””, etc. At a “phenomenal” level, in terms of its uses, purposes and symbols, language is diverse and non-uniform. But “metaphysically”, at the level of the relationship between language and reality, Lovibond claims that all language stands on the same level, contrary to the “heterogeneous” metaphysical picture on which there is language which enjoys a privileged metaphysical relationship with reality and language which does not. If Lovibond was asserting that language is homogeneous at the “phenomenal” level, Diamond’s criticism would hit its mark. But since Lovibond is careful not to make that claim, I do not think Diamond’s objection works.

Whether Lovibond is right that we should endorse the metaphysically homogeneous view of language is a question I will not investigate in this thesis, since I believe I can successfully undermine the Fact/Value Dichotomy without appealing to Lovibond's metaphysically homogeneous view of language (see Chapter 3, Sections 1 and 2, below). For the same reason, I do not need to spend too much time probing Lovibond's (possibly controversial) distinction between "metaphysical" and "phenomenal" claims about language. I am happy to grant Lovibond that a metaphysically homogeneous view of language is preferable to a metaphysically heterogeneous view, according to the basic understandings of those terms I gave above. I shall now argue that, even if Lovibond is right to adopt the metaphysically homogeneous view of language, her argument against the Fact/Value Dichotomy is still unsuccessful.

Part ii: Problems with Lovibond's Argument

Lovibond's strategy is to try to argue that the Fact/Value Dichotomy depends on the metaphysically heterogeneous view of language, but that since, she says, that view should be rejected in favour of the metaphysically homogeneous view, we should also reject the Fact/Value Dichotomy. But even if we reject the metaphysically heterogeneous view, must we deny the Fact/Value Dichotomy?

Denying the metaphysically heterogeneous view certainly cuts off *one* route to the Fact/Value Dichotomy, namely that which uses the metaphysically heterogeneous view to motivate the idea that "thin" descriptive language and "thin" evaluative language are of two fundamentally different semantic kinds. But does the Fact/Value Dichotomy *have* to be based on the metaphysically heterogeneous view? I think not.

Consider Parfit (2011), who endorses the Dichotomous Claim, but does not appeal to a metaphysically heterogeneous view of language to support it. Rather, for Parfit, the Dichotomous Claim is supported by a *prima facie*, supposedly pre-theoretical inspection of ordinary language. It is simply obvious to a competent speaker, Parfit suggests, that words

like “right”, “ought” and “duty” are of a different kind to words like “red”, “square” and “tall” (2011: 265). So, Parfit may say, in reply to Lovibond: the Dichotomous Claim seems to fit with ordinary language, for it seems competent speakers recognise some sort of difference in kind between “thin” descriptive and “thin” evaluative words, and perhaps (Parfit and other advocates of the Dichotomous Claim will suggest) the most natural way of explaining this is to suppose that words in one of those categories have one kind of meaning, and the words in the other category have another.

This is an especially intractable problem for Lovibond, because she aims to support her own, different view by appeal to ordinary language, too. But, as Parfit’s defence of the Dichotomous Claim indicates, it is not clear either way whether ordinary language is consistent with the Dichotomous Claim or not. Parfit gives what seems like a plausible defence of the Dichotomous Claim by appeal to ordinary language. Does Lovibond’s view fare any better than Parfit’s view when it comes to its compatibility with ordinary language? If it does not, then, I would argue, that constitutes a serious problem with Lovibond’s argument, because then Lovibond will be forced to accept that the Fact/Value Dichotomy is still viable, so long as one does not try to support it through the metaphysically heterogeneous view of language.

To work out how Lovibond’s view fares with respect to ordinary language, we need to consider how she spells out her alternative to the Fact/Value Dichotomy. On the basis of her rejection of the “empiricist” metaphysically heterogeneous view of language, Lovibond claims that the attempt to “drive a wedge” between “descriptive meaning” and “evaluative meaning” fails, and therefore “[the descriptive] function pervade[s] all regions of discourse irrespective of content” (1983: 27). Even more strongly, she continues:

What, then, becomes of the other function of language recognized by empiricism – the “expressive” function which corresponds to the “value” term of the fact/value distinction?

...

The answer is that on the view we are now considering, this function also comes to pervade language in its entirety. Linguistic meaning as such takes on the character identified by the

empiricist as “emotive” or “expressive”. In other words, “fact” and “value”... *coalesce* – and assertoric discourse is now seen to accommodate both impartially. (1983: 27, emphasis in original)

What Lovibond suggests here, I believe, is a view on which *all* meaningful language is both descriptive and evaluative (i.e. semantically thick). She is not simply saying, however, that she believes all language *as a matter of fact* is “thick”, as it were, contingently. As we saw in Chapter 2, Section 1, the Dichotomous Claim is consistent with the possibility that all current language is, contingently, thick, because the Dichotomous Claim merely says that thin language is *possible*. But Lovibond’s view is different, because she writes “linguistic meaning *as such*” (emphasis mine) has the feature that the empiricist takes to be characteristic of “emotive” or “expressive” language. It seems to me that Lovibond is suggesting that all meaningful claims are *necessarily* both descriptive and evaluative. So, on her view, description and evaluation “pervade language in its entirety” not contingently, but rather essentially.

Such a view is certainly different from the Fact/Value Dichotomy. Of course, an advocate of the Humean Claim would deny Lovibond’s view that meaningful language is necessarily both descriptive and evaluative, and that assertoric discourse “accommodates” both description and evaluation, since the Humean only allows for semantically “thin” descriptive and evaluative claims. Moreover, Lovibond’s view is inconsistent with the Dichotomous Claim, since (on my interpretation of her view, at least) she denies the possibility of semantically “thin” descriptive or evaluative claims, since she claims that “linguistic meaning as such” necessarily includes both “descriptive meaning” and “evaluative meaning”.

Lovibond has, then, successfully articulated an alternative to the Fact/Value Dichotomy. However, I believe her view faces two major problems. First, it is plausible, at least *prima facie*, that the Dichotomous Claim is actually better-supported by ordinary language considerations than Lovibond’s view. At the very least, we shall see, it is hard to rule out this possibility. Secondly, Lovibond’s view seems to involve a kind of essentialist claim about linguistic meaning (a claim about linguistic meaning “as such”), which I believe is a

very strong claim that is hard to defend. I will now address each of these problems for Lovibond's view in turn.

Lovibond claims that the “descriptive function” and the “evaluative function”⁹⁴ “pervade all regions of discourse irrespective of content” (1983: 27). Is this plausible, though? When we examine ordinary language, *does* it all appear to be “thick”? Most of us would, I think, answer “no” (cf. Parfit 2011: 265). On the face of it, the words Parfit mentions, like “right”, “ought” and “red”, *do* appear to be “thin”, i.e. “purely” descriptive or purely evaluative. So, for that matter, do many claims which feature such words, such as “your action was right”, and “the square is red”.

The reason this is a serious problem for Lovibond's view is that Lovibond's main reason for rejecting metaphysically heterogeneous views of language, and hence the Fact/Value Dichotomy, was based on ordinary language considerations. As she puts it, “philosophical considerations cannot discredit the way in which we classify linguistic entities for other, non-philosophical purposes” (1983: 26), and she thinks non-cognitivists fall foul of this rule, with their revisionary willingness to classify moral claims as non-propositional despite their propositional appearance.

As I have explained, though, it is possible to endorse the Fact/Value Dichotomy for reasons unrelated to the metaphysically heterogeneous view of language, as Parfit does. For Parfit (2011), the Dichotomous Claim simply fits the linguistic data best. So while Lovibond *may* be correct to reject metaphysically heterogeneous views of language, my point is that her appeal to ordinary language and her rejection of the metaphysically heterogeneous view are not enough to establish her view, since she also needs to reckon with Parfit's different way of defending the Fact/Value Dichotomy (which is, much like Lovibond's view, dependent on the “appearances” of ordinary language, see Chapter 2,

⁹⁴ It is not entirely clear to me why Lovibond uses the word “function” in this context, when in other places she talks about descriptive and evaluative “meaning”. To be charitable to Lovibond, I will assume that she is using the word “function” to mean the same as “meaning”, and that she is therefore claiming that all language is necessarily “thick” at the semantic level (which, as I have discussed, would make her view a genuine alternative to the Fact/Value Dichotomy).

Section 1). Since Parfit's view seems to be *at least as*, if not *more*, supported by ordinary language considerations than Lovibond's view, perhaps we should conclude that Lovibond's argument is not a successful rebuttal of the Fact/Value Dichotomy.

It is controversial, however, whether Parfit's view is as readily compatible with ordinary language as it might seem. Chappell (2013), for example, argues that we should *not* understand terms such as "good" and "right" as thin. While conceding that such terms can appear to be thin, Chappell (2013) gives an example to try to persuade us that these appearances are deceptive. According to Chappell (2013: 187), Japanese claims featuring the word "*giri*" are conventionally translated into claims in English featuring one of the words "right" or "wrong" (i.e. seemingly thin evaluative words). Yet if we probe a little deeper, we find features of the word "*giri*" which make it look more like a thick term: it involves, for example the specific notions of debts of gratitude and the idea that an agent who acts in a way that is "*giri*" is self-sacrificing in the pursuit of happiness (Chappell 2013: 188). So as Chappell says, "*giri*" "has a whole variety of historically particular and socially situated implications and applications, which together enable us to say a great deal about the factual content of "*giri*" (Chappell 2013: 188). In other words, saying that somebody acts in a way that is, or fits with, "*giri*" carries at least some factual descriptive content, as well as evaluative content.

The reason Chappell brings up this example is to motivate the idea that even concepts which appear "thin", at first glance, often turn out to be thick upon closer analysis. While "*giri*" is just a single case, and we should therefore not draw any sweeping conclusions based on that example alone, Chappell thinks the point does generalise. To support this, Chappell considers the historical and cultural specificity of seemingly "thin" moral concepts in English, "good" and "right", pointing out that, according to Anscombe (1958), those concepts are "extremely particular... historical products" (Chappell 2013: 189). The fact that we tend to *think* of them as generic and "thin" concepts is to be explained by our "lack of analytical distance" from them, Chappell suggests (2013: 188).

Rather than analysing the arguments of Chappell and Anscombe in detail here, I will remain neutral, in this thesis, on whether concepts like “good” and “right” are thick. For reasons I will state presently, there is no need for me to intervene on this issue.

While Parfit is, I think, certainly right that they *appear* to be thin at first glance, Chappell may be right that this appearance does not survive further scrutiny. Further argumentation would be needed to support Chappell’s view (particularly on whether apparently thin *descriptive* concepts like “red” are genuinely thin, on which see Chapter 3, Section 2, below), but if it can be sustained, this would seem like good news for Lovibond’s view. It would also be good news for my view (to be presented in Chapter 3), since I ultimately also wish to deny Parfit’s argument for the Dichotomous Claim. But my argument against the Fact/Value Dichotomy does not depend on Chappell being right on this count.

The apparent tension between Lovibond’s view and ordinary language *may*, then, dissolve (depending on whether Chappell and Anscombe are right that even words like “good” and “right” are thick). The reason I am nevertheless happy to remain neutral about Chappell’s view that concepts such as “good” and “right” are thick is that *the Dichotomous Claim is compatible with that view*. As I explained in Section 1 of this Chapter, it is consistent with the Dichotomous Claim to say that all present language is in fact thick. The Dichotomous Claim merely claims that semantically thin descriptive and evaluative language is *possible*, not that it is actual. To properly assess the Dichotomous Claim, then, we must move beyond simply analysing our ordinary language. We need to scrutinise the claims about which kinds of language are *possible*: specifically, both Lovibond’s view that semantically thin language is *impossible* on account of all meaningful language being both descriptive and evaluative “as such”, and the Dichotomous Claim that semantically thin language is possible.

As well as the problem about her view’s support (or lack thereof) from ordinary language, I believe Lovibond’s view faces another serious problem, severely limiting its effectiveness against the Fact/Value Dichotomy. This second, and ultimately more decisive, problem is that it appears to depend on a kind of *essentialism* about meaning.

Moreover, I do not believe Lovibond offers convincing reasons to accept the form of essentialism on which her view depends.

Recall that, for Lovibond, meaningful language “as such” is both descriptive and evaluative. Even if we accept that all present ordinary language is thick, though, what reason does Lovibond offer to think that all *possible* language is thick? Why could there not be semantically thin descriptive or evaluative language?

As we saw above, Lovibond tries to motivate her position by undermining the metaphysically heterogeneous view of language. Again, let us grant, for the sake of argument, that Lovibond is right to reject that view. For the rejection of the metaphysically heterogeneous view to take us all the way to Lovibond’s conclusion that there can be no semantically thin language, however, Lovibond would need to argue that *there can only be semantically thin descriptive and/or evaluative language if the metaphysically heterogeneous view is true* (i.e. if some language has a privileged metaphysical relation to reality and some does not). Not only is an argument for this premise absent from Lovibond’s (1983) book, I am sceptical that the premise is true. I can see no reason why the Dichotomous Claim would require the truth of the metaphysically heterogeneous view of language. All the Dichotomous Claim needs, after all, is for there to be two different kinds of possible meaning, one of which can be possessed by “thin” descriptive language, and one of which can be possessed by “thin” evaluative language. It is entirely unclear why that view would necessarily be threatened by an argument against the metaphysically heterogeneous view of language. For example, Parfit’s (2011) defence of the Dichotomous Claim, as considered in Chapter 2, Section 1, appeals to ordinary language considerations rather than any metaphysical thesis.

Chappell (2013), differing from Lovibond, readily admits that even if “good” and “right” are thick concepts, this does not in itself show that there cannot be thin concepts.⁹⁵ So the

⁹⁵ This shows, incidentally, that Chappell’s view is not, and indeed is not intended to be, threatening to the Dichotomous Claim, which is why it is not necessary for me to decide whether Chappell is right.

only way of motivating Lovibond's radical thesis is to make plausible the idea that meaningful language is *essentially* thick. As I have just argued, Lovibond's rebuttal of the metaphysically heterogeneous view of language does not do this. Moreover, it is hard to imagine any other argument which *could* make this very strong conclusion plausible. While concepts like "good" and "right" may, as they are typically used in English, after all be thick, as Chappell claims, what is the contradiction in supposing that we might be able to coin a meaningful thin evaluative word or phrase, constituted by a watered-down, non-descriptive notion of rightness, for example as a generic term of positive commendation? Lovibond's dismissal of the metaphysically heterogeneous view, and her appeal to ordinary language, do not, as far as I can see, yield reasons for thinking that the notion of semantically thin language is contradictory or otherwise problematic.

My point here is that trying to defend Lovibond's view seems very arduous. For one thing, the Dichotomous Claim just seems very *prima facie* plausible, whereas Lovibond's essentialism about meaning will, I believe, seem very counterintuitive, especially to philosophers like Parfit (2011). Secondly, since Lovibond does not give sufficient argumentative support for this essentialist view, we would need to find some new arguments in favour of it. But the essentialist view of language Lovibond commits to is a very strong view, in the sense that it entails a lot. It entails, for example, not only that all actual language is semantically thick but that all *possible* language is thick, which is even harder to prove than the claim that all present language is thick. Crucially, on top of this, Lovibond's rejection of the metaphysically heterogeneous view of language does not refute the Parfit-style defence of the Dichotomous Claim by appeal to ordinary language.

If some argument which adequately defends Lovibond's view is found, then I would be happy to accept it, since I do not accept the Fact/Value Dichotomy. However, opponents of the Fact/Value Dichotomy may be forgiven for looking for an alternative strategy to Lovibond's: an alternative which places a less weighty burden on our shoulders. My conclusion on Lovibond's view, then, is not necessarily to say that it is false or otherwise wrong-headed, but that it is a very strong view which has not yet been adequately defended. Since, I believe, defending her view would be highly arduous, for the reasons I have just mentioned, I will be proposing and defending an alternative strategy for opposing

the Fact/Value Dichotomy in Chapter 3, which deliberately requires less strong claims than Lovibond's (for example, by avoiding the need for essentialism). This, I will argue, makes my argument against the Fact/Value Dichotomy significantly more plausible and successful than Lovibond's.

Part iii: Desiderata for an Alternative anti-Dichotomy Strategy

I have now considered two of the most influential arguments against the Fact/Value Dichotomy in the literature on this topic, those from Putnam (2002) and Lovibond (1983). Although I found both wanting, I think we can draw some useful conclusions and insights from our investigation of them. Specifically, I will now identify some desiderata that an improved argument against the Fact/Value Dichotomy would need to satisfy.

As I argued in Chapter 2, Section 2, the main problem with Putnam's argument is that it does not apply to the Dichotomous Claim, but rather the Humean Claim only. Putnam resists the idea that thick concepts can be separated into thin components, and denies that the "pictorial semantics" model accurately represents all of our descriptive language, but I argued that neither of these moves speak against the Dichotomous Claim's commitment to the possibility of thin language. Our first desideratum for an improved strategy is hence for it to explicitly target both the Dichotomous Claim and the Humean Claim, and to be effective against both. Of course, if we cannot find a good argument against the Dichotomous Claim, it would likely be more reasonable to accept the Dichotomous Claim (see Section 1 of this Chapter, above). But I believe there is a good argument against the Dichotomous Claim to be made, which I set out in Chapter 3.

Putnam and Lovibond both target certain assumptions about language and/or meaning, which they believe have led philosophers to, erroneously, endorse the Fact/Value Dichotomy. For Putnam (2002), this is "pictorial semantics", whereas for Lovibond (1983) it is the metaphysically heterogeneous view of language. While I think Putnam and

Lovibond are right to critically target the ideas about meaning underpinning the Fact/Value Dichotomy, critics of the Fact/Value Dichotomy need to make up their minds about *which* exact ideas about language and meaning truly underpin the Dichotomy, and why they are mistaken. So my second desideratum for an improved strategy is for it to give a precise and clear statement of which (if any) false views about meaning genuinely underpin the Fact/Value Dichotomy. Is it really “pictorial semantics” which is responsible, or Lovibond’s notion of the “metaphysically heterogeneous” view of language? Drawing from Kovesi (1967), I will argue that neither Putnam nor Lovibond quite lay their finger on the right point here, although there are recognisable similarities between my diagnosis of the problem with the Fact/Value Dichotomy and theirs (see Chapter 3, Section 2).

Finally, as suggested above, I want to offer a way of resisting the Fact/Value Dichotomy which does not depend on the radical essentialist claims needed to reinforce Lovibond’s position. To repeat, if Lovibond’s position *can* be defended adequately, I will be cheering the demise of the Fact/Value Dichotomy along with her. In the absence of such a defence, however, opponents of the Fact/Value Dichotomy would do well to make an argument against the Dichotomy which depends on less ambitious claims. Specifically, I want to make do without appeal to any sort of essentialism about language or meaning, which seems to me the most difficult part of Lovibond’s argument to accept or justify. By providing such an alternative argument, I believe I will, therefore, make a significant improvement upon previous attempts to oppose the Fact/Value Dichotomy. Even if the reader judges my subsequent argument to ultimately fail, then, I believe there is some value in identifying an anti-Dichotomy strategy which is less onerous, and therefore more plausible, than Lovibond’s: although my specific attempt to make good on the strategy may fail, outlining the less onerous strategy can help philosophers recognise that the Fact/Value Dichotomy is less safe from criticism than it may appear, and potentially develop other more successful arguments through the route I identify.

Conclusion to Chapter 2, Section 3

Lovibond attempts to undermine the Fact/Value Dichotomy by undermining the “metaphysically heterogeneous” view of language, which she believes the Dichotomy to depend on. Instead of embracing the Dichotomy, she instead argues that meaningful language “as such” is semantically thick. If true, this means there cannot be thin descriptive or evaluative claims, so both the Humean and Dichotomous Claims are false.

After setting out Lovibond’s case against the Fact/Value Dichotomy, though, I drew attention to three main problems with her argument. First, even if the metaphysically heterogeneous view of language should be rejected, as Lovibond argues, this does not rule out the Fact/Value Dichotomy, as Parfit’s (2011) defence of the Dichotomous Claim shows. This limits the effectiveness of Lovibond’s strategy. Second, it is controversial whether Lovibond’s view is readily compatible with ordinary language, as she believes it is. In fact, Parfit’s considerations regarding the Dichotomous Claim seem to indicate that the Dichotomous Claim may fit better with ordinary language than Lovibond’s view. However, since Chappell (2013) successfully casts doubt on this, I do not consider this problem fatal for Lovibond’s argument. More to the point is the essentialism about meaning that Lovibond needs to commit to for her argument to work. This, I have argued, is very difficult to defend, and Lovibond does not offer enough in defence of it. Her attack on the “metaphysically heterogeneous” view of language, I argued, does not support the essentialism she needs, since one can reject the metaphysically heterogeneous view without adopting her essentialism.

The conclusion I draw from this is that critics of the Fact/Value Dichotomy would be in a much stronger dialectical position if we can find some way of opposing the Dichotomy without resorting to an essentialism of the kind that Lovibond sets out. If that is the only way to resist the Dichotomy, then the Dichotomy (and, I think, the Dichotomous Claim, specifically) will look more plausible than ever. My most important desideratum for my alternative argument against the Dichotomy is hence for it to reduce the plausibility of the Fact/Value Dichotomy *without* committing to any kind of essentialism about language or

meaning, as Lovibond does. Fortunately, as we will soon see, I do believe that critics of the Dichotomy can motivate their position without appealing to essentialism of any kind. To do this, we will need to draw from, and develop, ideas from Kovesi (1967).

To recap, and to make it easier to refer back to them later, I will now list my three desiderata for an alternative strategy for resisting the Fact/Value Dichotomy:

- 1) *Completeness* Desideratum: For our argument to tackle both the Humean and Dichotomous Claims,
- 2) *Linguistic* Desideratum: For our argument to pin down exactly which incorrect views or assumptions (if any) about language underpin that Dichotomy, and
- 3) *Non-Essentialism* Desideratum: For our argument not to rely on an essentialism about language or meaning, as Lovibond's does.

If an alternative approach can deliver all these desiderata, I believe it would mark a significant improvement upon the arguments from Putnam and Lovibond, in terms of how plausible and effective it would be at resisting the Fact/Value Dichotomy. I believe the argument I will present in Chapter 3 achieves this, as I will argue below.

Conclusion to Chapter 2

I will now briefly summarise Chapter 2. First, I distinguished between the Humean Claim and Dichotomous Claim: a stronger and weaker form, respectively, of the Fact/Value Dichotomy. The Humean Claim says that all descriptive and evaluative language is semantically “thin”, and that descriptive claims and evaluative claims have different kinds of meaning from one another. The Dichotomous Claim, by contrast, allows that there can be semantically “thick” (both descriptive and evaluative) language, but also allows the possibility of semantically “thin” language, and says that semantically “thin” descriptive language has (or would have) a different kind of meaning to semantically “thin” evaluative language.

I argued that while Johnston defends the stronger version, the weaker version is more plausible, and is fairly well-supported by Parfit's (2011) ordinary-language considerations. I considered these two versions of the Dichotomy in relation to the literature on "thick" concepts, and found that while both versions of the Dichotomy seem to have good ways of explaining the phenomenon of apparently-thick concepts, the Dichotomous Claim has an easier time doing so. However, by expressing some worries about standard "Inseparabilist" views on thick concepts, I motivated further critical scrutiny of the Dichotomous Claim as well as the Humean Claim.

I then explored the arguments of two critics of the Fact/Value Dichotomy: Putnam (2002) (in Section 2) and Lovibond (1983) (in this Section), but found that neither of them are successful against the Fact/Value Dichotomy as I define it. Putnam's view only targets the Humean Claim, while Lovibond's view depends on a very strong form of essentialism. Accordingly, I drew out three desiderata for an improved argument against the Fact/Value Dichotomy, which I will use to generate and support my argument in Chapter 3.

Chapter 3, Section 1: How to Resist the Fact/Value Dichotomy without Essentialism

The Dichotomous Claim, let us recall,⁹⁶ is that it is possible for there to be semantically “thin” descriptive and evaluative language. In the previous Section, we considered Lovibond’s attempt to refute this claim by advancing a form of essentialism. If all meaningful language “as such” is essentially semantically thick, then the Dichotomous Claim would be false. I have argued, however, that opponents of the Dichotomous Claim (and the Humean Claim) would do better to advance their criticisms without relying on any kind of essentialism. In this Section, I set out a strategy for achieving this.

I am going to explain my strategy by considering how a very similar strategy can be applied to a different philosophical issue. The parallel case I will consider is Wittgenstein’s investigation of whether there could be a logically private language, focusing especially on Mulhall’s (2007) interpretation of Wittgenstein on this issue. Let me stress at the outset that none of the claims or arguments of this thesis depend on the truth of any of Wittgenstein’s or Mulhall’s claims about the possibility of private language. All I am doing is using the ideas of those two philosophers on the issue of the possibility of private language as an instructive foil against which my own argumentative strategy can be better understood.

In section 243 of *Philosophical Investigations* (Wittgenstein 1953), Wittgenstein asks:

Is it... conceivable that there be a language in which a person could write down or give voice to his inner experiences – his feelings, moods, and so on – for his own use? – Well, can’t we do so in our ordinary language? – But that is not what I mean. The words of this language are to refer to what only the speaker can know – to his immediate private sensations. So another person cannot understand the language.

⁹⁶ See Chapter 2, Section 1 for a more detailed explanation.

Wittgenstein, then, is asking whether a certain kind of language is possible. Specifically, he asks whether a “logically private” language is possible: one which nobody apart from the speaker can understand. We should note the modality of the notion “can understand” here: what is under consideration is not whether there can be a language which, contingently, only one person can understand. In this case, other persons *could* learn and come to understand the language, even if only one person understands it at present. Rather, Wittgenstein is considering whether there could be a language which is *in principle impossible* for anybody other than the solitary user to understand.

Much like those of us critically evaluating the Dichotomous Claim, then, Wittgenstein is here considering whether a certain kind of language is *possible*: in his case, a “logically private” language of the kind just explained. In our case, we are concerned with whether a semantically thin descriptive language is possible. As I discussed in Chapter 2, the contingent matter of whether any *present* ordinary language happens to be semantically thin is orthogonal to the truth of the Dichotomous Claim, since semantically thin language may be possible even if, contingently, we lack it. In much the same way, one could argue that a logically private language is possible even if we currently lack any such language. Since Wittgenstein is also, much like the present author, considering the possibility of a certain kind of language, it makes sense to explore whether the argumentative strategy he deploys on the issue of logically private language might be applicable (in a suitably adapted form) to the Dichotomous Claim.

We can imagine that one strategy for tackling the idea that a logically private language is possible would be to argue that language is essentially public. This would be, like Lovibond (1983) (see Chapter 2, Section 3, above), to appeal to a form of *essentialism* about language in order to show that a certain kind of language is not possible. If all possible language is *essentially* semantically thick, then there can be no semantically thin language. So if we can argue that all possible language is essentially semantically thick, we will have grounds for rejecting the Dichotomous Claim.

Mulhall (2007), however, suggests that this was not Wittgenstein's strategy in his attack on the idea of private language, and indeed it is not the argumentative strategy of the present thesis.

Instead, on Mulhall's interpretation, Wittgenstein challenges the idea that there can be a logically private language by attempting (so to speak, *on behalf* of the person claiming that there can be such a language, the would-be "private linguist") to *make sense* of that idea.

Mulhall explains:

Wittgenstein tries to imagine, and then tries out, ways of giving meaning to the constituent terms of the interlocutor's formation ... the succeeding investigations must determine whether there is any way of taking them that might give them genuine substance. In so far as Wittgenstein's attempts to imagine such ways fail to satisfy his interlocutor... then what the interlocutor does mean remains undetermined... It remains open to him [the interlocutor] to imagine another such way, and thereby to find the satisfaction he seeks; but if he does not, then Wittgenstein implicitly invites him to ask himself why he is passionately convinced that his words mean something in particular (Mulhall 2007: 18-19).

As I understand this approach, the idea is that, while it is clearly possible to string together the words "a language that nobody apart from the speaker can possibly understand" (or, for that matter, a "semantically thin descriptive language"), it remains an open question whether such a string of words is *meaningful* or *substantial* and hence designates a genuine possibility, rather than being empty.⁹⁷ To work out whether they are meaningful requires us to investigate and explore ways in which they might be meaningful. If we cannot find any such way of making them meaningful – of showing them to be substantial, rather than empty strings of words - that is not necessarily to show that no sense *could ever* be made of such words, since perhaps we simply did not find the right way of making sense of them. However, it would allow us to say that *until* or *unless* some way is found to make sense of these words, the words are, so to speak, empty – no sense has yet been made of them, and

⁹⁷ For related issues regarding the interpretation of Wittgenstein's *Tractatus*, see e.g. Diamond 2000 and Mulhall 2007, Ch. 1. For another interpretation of Wittgenstein on private language which is similar to Mulhall's, see Canfield (2001).

it is, at best, unclear whether they designate a genuine possibility. As Canfield (2001: 378) highlights, this Wittgensteinian method involves examining a set of philosophical statements and probing to see whether they are, after all, “disguised nonsense”, in which case it will be our job to “make [it] plain” that it is nonsense and not sense.

This approach has several attractive features that make it optimal for the present attempt at finding a flaw with the Fact/Value Dichotomy. First, it does not attempt to definitively *prove* that there can be no such thing as semantically “thin” descriptive (or evaluative) language, for instance by appealing to a kind of essentialism, as Lovibond does. This means that we can potentially make a more convincing case against the Dichotomy by making our argument depend on less strong claims than Lovibond’s. Hence, this approach satisfies the *Non-Essentialism* desideratum identified at the end of Chapter 2. But while I do not attempt to definitively disprove the Dichotomous Claim, I still think my argument is effective in undermining it (see Section 2 of this Chapter, below).

The second, related advantage of this view is that it is fair, and conduces to a more constructive engagement with the Dichotomous Claim than Lovibond’s approach. The reason is that, in executing this alternative strategy, we will need to constructively explore ways of making the Dichotomous Claim intelligible – to lay bare the logic behind the view. Much like with Wittgenstein’s investigation of the notion of private language, it will only be by genuinely attempting to make sense of the notion of thin language that we will be able to determine whether or not the advocate of the Dichotomous Claim ultimately succeeds in doing so. While I ultimately cannot find a way of making the notion of thin descriptive language intelligible, I leave it open whether such a way could be found in future.

Finally, while this approach is therefore attractively undogmatic, allowing that we might, in future, find some intelligible explanation of the notion of semantically “thin” descriptive or evaluative language, even if we cannot find one presently, it is still effective. If we cannot presently find any suitable explanation of the notion, the question would be raised why we should agree that a semantically thin language *is* a genuine possibility, after all.

The burden will then rest on the defender of the Dichotomous Claim to convince us that semantically “thin” language is possible, rather than *vice versa*. For Parfit, the “default” position is to accept the possibility of both thin and thick language, because he thinks that we can find examples of both through a cursory look at ordinary language. Against this, I shall argue that the “default” should be not to accept the possibility of thin descriptive language, unless or until the notion of such language can be satisfactorily explained.

By critically targeting the idea that a semantically “thin” descriptive language is possible, I am attacking both the Humean Claim and the Dichotomous Claim. I thereby satisfy the *Completeness* desideratum from Chapter 2, Section 3, and improve upon Putnam’s (2002) account, on which see Chapter 2, Section 2. The latter of these Claims openly depends upon the possibility of such language (see Chapter 2, Section 1), but the former does, too: since the Humean Claim is that *all* descriptive language is “thin” and *all* evaluative language is “thin”, it is equally tied to the possibility of semantically “thin” descriptive language as the Dichotomous Claim.

Finally, this approach lets us meet the *Linguistic* desideratum from Chapter 2, Section 3. In the next Section, I will offer my own account of which assumptions about language underpin the Fact/Value Dichotomy (which is different to the accounts of Putnam and Lovibond), and why these assumptions are mistaken. The approach to understanding “thick” language bound up with these assumptions seems to vindicate the Fact/Value Dichotomy, but I then argue that it is mistaken, drawing from Kovesi (1967).

The strategy suggested by Wittgenstein and Mulhall therefore seems to be apt for the issues we are dealing with in this thesis, namely the issue of the Fact/Value Dichotomy. I will therefore attempt to implement this strategy in the following Sections, to undermine the Fact/Value Dichotomy. By doing so, I meet my desiderata for an improved case against the Fact/Value Dichotomy and, I believe, put the Fact/Value Dichotomy in considerable jeopardy.

Chapter 3, Section 2: Why We Should Reject the Fact/Value Dichotomy

In the previous Section, I presented my strategy for undermining the Fact/Value Dichotomy. This is to challenge the advocate of the Dichotomy to make sense of the notion of semantically “thin” descriptive language. Now, I aim to execute that strategy by drawing on the work of Kovesi (1967). With Kovesi’s insights, I argue, we will see why advocates of the Fact/Value Dichotomy have not, as yet, successfully made sense of the idea of semantically “thin” descriptive language, on which their Dichotomy depends.

Recently, there has been renewed critical interest in Kovesi’s work (see Tapper and Mooney 2012). However, none of this work follows the strategy that I am opting for here, and, to my mind, none of it successfully repudiates the Fact/Value Dichotomy. Harrison (2012), for example, thinks that Kovesi appeals to the essential publicity of language to refute the Dichotomy. One reason I do not follow Harrison here is that I want to avoid any essentialist claim about language or meaning, as per the desideratum set out Chapter 2, Section 3, above. Secondly, as will become evident in what follows, my interpretation of Kovesi places no weight on the notion of the publicity of language, so I do not think we need to make any claims about that to understand the crucial flaws with the Fact/Value Dichotomy.

In response to the suggestion that the advocate of the Fact/Value Dichotomy may not be able to make *sense* of the notion of semantically “thin” descriptive language, Parfit (2011) might reply in the following way: “Of *course* it is possible to make sense of this notion! It is not only a coherent notion, it is a real phenomenon!” As we saw in Chapter 2, Section 1, Parfit maintains that words such as “red” and “square” are “thin” descriptive concepts, which can be used to formulate “thin” descriptive claims, like “the chair is red” or “the room is square”. He also maintains that words such as “good” and “right” are “thin” evaluative words, which can also be used to formulate thin evaluative claims. For Parfit, then, the notion of semantically “thin” language is perfectly intelligible, and readily

exemplified in our ordinary language.

If my strategy against the Fact/Value Dichotomy has any hope of success, then, I must be able to rebut Parfit on this score. As we will see, Kovesi's (1967) ideas about the meanings of words show that the matter is more complicated than Parfit realises. It turns out not to be straightforward to analyse the meanings of terms such as "red" and "square" in such a way that they turn out to be semantically "thin". Indeed, Kovesi, I shall argue, exposes that the reason we get the impression that these terms are "thin" is that we have a mistaken idea of what is required to know the meaning of a word. Kovesi hence gives us resources for objecting to Parfit's way of carving up descriptive and evaluative language, which, as I explained in Chapter 2, Section 1, part i), is crucial if opponents of the Fact/Value Dichotomy are to have any chance of success. Using Kovesi's arguments as a springboard, I will argue, contrary to Parfit, that a satisfactory explanation of the notion of semantically "thin" descriptive language is not yet forthcoming.

My argument below presses the likes of Parfit to explain *what it is* in virtue of which "thin" descriptive words are (supposedly) different in kind from "thick" words and "thin" evaluative words. I find that there is no convincing rationale motivating the idea that there is a difference in kind here, and this exposes that the notion of semantically "thin" descriptive language has not been adequately explained.⁹⁸ This gives us reason to provisionally reject the Dichotomous Claim, in line with the strategy presented in Section 1 of this Chapter.

Part i: Kovesi on Material and Formal Elements

I will now sum up the relevant ideas from Kovesi (1967) which will help support my

⁹⁸ My argument focuses on semantically thin descriptive language, and I will not pass judgment on whether there can be semantically thin evaluative language in this thesis. However, since the Humean and Dichotomous Claims both commit to the possibility of semantically thin descriptive language as well as semantically thin evaluative language, targeting the notion of semantically thin descriptive language is enough to make an effective argument against both forms of the Fact/Value Dichotomy.

argument against the Fact/Value Dichotomy.

Why focus on Kovesi's views here, since the debate on this issue has progressed so much since Kovesi was writing? The simple reason is that most of the subsequent debate on the relationship between descriptive and evaluative language⁹⁹ (e.g. the debate on "Separabilism" vs "Inseparabilism" about thick concepts, on which see Chapter 2, Section 1, parts iii)-iv)) seems to take the Dichotomous Claim for granted, which I think is a mistake. Even Chappell (2013), who is highly critical of the notion of "thin" language, still seems to grant that semantically thin concepts are possible, as I mentioned in Chapter 2, Section 3, part ii). Hence, I return to Kovesi because I think his approach, unlike those of participants in the more recent debates, helps us see the serious flaws with the Dichotomous Claim (in this respect my approach is similar to that of Harrison 2012).

Kovesi insists on a distinction between "formal elements" and "material elements" in connection with the meaning of a word, and this distinction is crucial to his resistance to the Dichotomous Claim. So my first task is to survey what Kovesi says about this supposed distinction.

Kovesi introduces this distinction with the example of the word "table". He defines his notion "material element" by saying:

I call the *matter* [or "material element"] not only the various materials out of which we may construct tables but any characteristics in which the object may vary without ceasing to be a table. So the shape of the table, far from being its *form* will be part of its matter, since it may change - a table may be oblong, round or square - while the object remains a table. Similarly, whether the object has four legs or three, whether it is made of wood or iron, are questions about the material elements of tables. (1967: 4)

⁹⁹ With the exception of Lovibond (1983), whose views I have already discussed and found insufficient (see Chapter 2, Section 3).

First, a couple of clarificatory points about what Kovesi's notion of "material element" is, and is *not*, supposed to be. First, it is *not* supposed to apply only to words which denote tangible physical objects, Kovesi says (1967: 4). It can apply equally well to words like "murder": "it is the material element of an act of murder that someone drives a knife into his victim's heart, or administers poison, or strangles him", etc. (1967: 4). These are examples of material elements because they are examples of *characteristics of the thing in question that could change without the thing (in this case, murder) ceasing to be what it is (i.e. a murder)*. So, for instance, if we realise it was a sword, not a knife, which was used to do the deed, we have a change in the material elements of the murder, but not in the fact that it was a murder (1967: 4).

Secondly, nothing in Kovesi's definition commits him to the idea that "material elements" must be observable properties. A material element is simply any "characteristic" which could change without the thing ceasing to be what it is, i.e., it cannot be an essential characteristic.¹⁰⁰ Nevertheless, in his examples, Kovesi always seems to give observable properties as examples of material elements. So, typically at least, it seems that by "material element", Kovesi means the non-essential observable properties of the thing in question.

What, then, of Kovesi's notion of "formal element"? Kovesi explains that "an answer to the question why we call a large variety of objects "tables" and refuse the word to other objects gives what I want to call the *form* of a table" (1967: 4). Unfortunately, this explanation is rather opaque. To understand what Kovesi means, we need to consider his contrast between the word "table" and colour-words, such as "yellow".

For Kovesi, the salient difference between "table" and "yellow" (and other colour-words) is that, for the latter, there is a single "perceivable quality" that we perceive when we see something that looks yellow, i.e. the colour (which we call yellow). But although some

¹⁰⁰ There is, of course, a complex literature on the topic of essential characteristics, or properties: see e.g. Marcus (1967) and Kripke (1980). I will not adopt any particular controversial stance on these issues in this thesis.

philosophers may try to argue that, in the same way, there is a perceivable property of “tableness” which we perceive when we see something that looks like a table, Kovesi thinks this is implausible. The reality, Kovesi says, is that “there is no perceivable quality called “tableness” as there is a perceivable quality called “yellow”” (1967: 2-3). It remains true, however, that there are certain physical properties that are *characteristic* of tables: for instance, a flat surface, legs, etc (1967: 3). Kovesi says “certain qualities must be present in a piece of furniture in order that we should be able to call it a “*table*”, but there is no strict rule as to what these qualities must be... on the other hand not just anything will qualify as a table” (1967: 3, emphasis in original).

The upshot is that the case of “yellow” is in a certain respect much simpler than that of “table”. We know which quality must be present in something for it to be appropriate to call it “yellow”: simply the colour we call “yellow”. In the case of “table”, by contrast, while there are several qualities that we *typically* associate with tables, like those Kovesi lists, the rule here is not so straightforward. Knowing about the qualities we typically associate with tables is clearly relevant to understanding the word and being able to use the word correctly. If some person did not know that tables typically have legs, a flat surface, and are about waist height, we would rightly doubt whether they know the meaning of “table”. Yet since even competent speakers cannot give a final list of qualities that would be necessary and/or sufficient for something to be a table, Kovesi thinks that being competent in using the word “table” – and understanding that word – involves and requires more than just awareness of the typical physical qualities of tables.

What is this extra ingredient? Kovesi writes that “Historically, without a need for tables we would not have these pieces of furniture; logically, we cannot understand the notion of a table without understanding that need” (1967: 3). So, for Kovesi, we cannot understand the notion of a table without understanding our need for tables, that is, the *purpose* that we expect tables to serve. We have a need for places to keep cutlery at hand while eating, and we have a need for surfaces on which to write at a convenient height, etc. (1967: 3). Correspondingly, the *purpose* of a table is to allow us to do these things (and potentially other things too).

Recall that we seemed to face a predicament about how to determine which objects count as tables. On the one hand, we know that the physical qualities of an object are relevant to whether we should call it a “table”. But at the same time, we cannot provide a definitive list of those qualities. Kovesi thinks we can clear up our puzzlement here by considering how our knowledge of the *purpose* of tables helps us distinguish between genuine tables and other objects. His idea is that we can work out whether a given object is a table based on whether the object in question is sufficiently well-suited to serve the purposes tables are supposed to serve; if so, a competent speaker will call it “table”.

Kovesi uses the example of the word “dirt” to illustrate this point. Somebody may perfectly well be able to tell us about the material elements of dirt, e.g. that it is typically a small collection of organic matter. Yet if they do not understand why custard on a shirt is dirt, but custard on a plate is not, or why sand on a lens is dirt, but sand on a beach is not, they do not yet have full grasp of the concept “dirt” (1967: 38). This sort of deficiency in understanding relates not to one’s awareness of the material but the formal element of the concept, since we reserve the word “dirt” for cases where small bits of typically-organic matter end up where they ought not to be, for e.g. aesthetic or prudential reasons. Part of the function of the concept is to draw attention to these undesirable instances, and somebody who failed to grasp that would thereby indicate that they do not fully understand the concept.

Part ii: Kovesi and the “Purpose Thesis”

Rather than dwelling on whether Kovesi’s explanation of the distinction between “material elements” and “formal elements” is satisfactory, or indeed on whether there are such things as formal or material elements, I want to focus on what we can learn from Kovesi’s discussion about what it is to understand the meaning of a word. With his terminology of “material” and “formal” elements, I believe Kovesi is trying to highlight that there are two distinct components involved with understanding the meaning of a word.

Consider first one's awareness of what Kovesi calls the "material element". For example, if one is aware that tables are typically waist height, have flat surfaces, have legs, etc., one has part of what one needs to correctly apply, and understand, the word "table". The evidence for this is that, as Kovesi points out, somebody who we found to be unaware of these typical properties of tables would rightly be judged not to fully understand the word "table". For example, if somebody claimed that tables do not typically have flat surfaces or legs, it would be more natural to conclude that they do not understand the meaning of "table" than that they know the meaning of the word, but are factually mistaken.

In further support of Kovesi's point that awareness of "material elements" (e.g. flat surfaces and legs, in the case of tables) is an important component of understanding the meaning of a word,¹⁰¹ consider that those who are aware of the typical "material elements" of tables will naturally be better at identifying which objects are tables, and this is crucial to mastering the word "table". Someone who is good at distinguishing tables from non-table objects is less likely to misapply the word "table" by calling something that is not a table a table. Conversely, a person (such as a child) who does not yet (fully) grasp the concept "table" may often make mistakes when asked to distinguish between objects that are, and are not, tables. The child improving their ability to distinguish tables from non-tables is part of developing their mastery of the use of the word "table". Since acquiring awareness of the typical "material elements" of tables is crucial to this process, Kovesi's idea that this is an important part of what it takes to understand the meaning of a word is plausible.

However, Kovesi also thinks there is another component involved with understanding a word, which concerns one's awareness of the "formal element". This concerns one's awareness of the relevant purpose or purposes of the word.¹⁰² In what way? To repeat

¹⁰¹As we will see below, the advocate of the Fact/Value Dichotomy takes this (correct) idea and misapplies it, by arguing that some words can be completely understood *exclusively* through recognition of the "material elements", and through being able to distinguish things which fall under the word and things which do not. Kovesi's point is that while this is *part* of what is involved with understand the meaning of a word, it is not all that is involved.

¹⁰²It is, of course, important to distinguish between the purpose of a word, and the purpose of an object or kind of object which a word refers to. Kovesi's point is simply that, often (but not always), the purpose of a

Kovesi's example, somebody may know what dirt typically looks like, yet fail (and perhaps systematically) to apply the word in the appropriate way, if they are unaware of the purpose of the word. Since we do not simply say "dirt" in response to any dirt-like appearances, but only in particular circumstances (e.g. when custard is on a shirt rather than a plate), a speaker needs to grasp the purposes associated with the word in order to be able to apply the word "dirt" correctly. We distinguish between correct and incorrect uses of the word "dirt" not only on the basis of whether the speaker has identified the right sort of object, but also based on whether their application fits the purposes that we expect that word to serve. I will refer to Kovesi's view that understanding a word requires awareness of the purpose of the word as the "purpose thesis".

As Queloz (2019) points out, the idea that words and concepts can be understood by considering their point or purpose has become increasingly influential in recent literature.¹⁰³ Queloz also points out, though, that the notion of the "point" of a word or concept is ambiguous. So, to be clear, the kind of "point" of a concept that Kovesi and I are focusing upon here is what Queloz usefully calls the "practical point": "the salient practical difference which the concept actually makes to the lives of concept-users" (2019: 1128). In Heal's (2013: 342) terms, to show the practical point of a concept is to show how it "enables or enriches realization of one or more of our interests".

To understand this idea better, consider the application of this approach from Craig (1990) as applied to the concept of knowledge. Craig explicates the concept of knowledge by considering the "practical point" (to use Queloz's phrase) of the concept: specifically, its role in helping those who use it to acquire information about their environment and in distinguishing between reliable and unreliable sources of information. Craig argues that understanding the concept in this light removes the need to analyse the concept in terms of necessary and sufficient conditions (Craig 1990: Chapters II-III, cf. Queloz 2019: 1122). While it is beyond the scope of this thesis to assess whether Craig's specific account is

kind of object will feed into the purpose of the word we use to refer to objects of that kind. For example, since we need tables to eat and write on, part of the purpose of the word "table" is to help us identify objects which are suitable for those purposes (1967: 3). In other kinds of case, the purposes of the word may be unrelated to the purpose of any particular object, such as when the word in question is not the word for a (kind of) object (e.g. "murder").

¹⁰³ Queloz's examples include Greco and Henderson (2015) and Fricker (2016).

plausible, his attempt helps us see how, as Kovesi urges, looking to the practical point of a concept seems crucial for fully understanding that concept.

With his distinction between “material” and “formal” elements, Kovesi, as I read him, is motivating a specific view about meaning, and what it takes to grasp the meaning of a word. Put simply, his idea is that we can fail to understand the meaning of a word in at least two different ways. One can manifest a failure to grasp the meaning of a word by lacking the necessary awareness of the qualities typically present in the things which that word refers to, as we saw with the example of the child who systematically fails to distinguish tables from non-tables correctly. Kovesi (1967: 41) calls these qualities “recognitors”, because they help competent speakers recognise the objects in question, and name or describe them appropriately. The child who is not aware of the relevant recognitors, and who hence systematically fails to distinguish tables from non-tables correctly, does not yet fully grasp the meaning of the word “table”.

Another distinct kind of failure is that displayed in Kovesi’s “dirt” example, of the person who does not understand why custard on a shirt is dirt, but custard on a plate is not. To fully grasp the meaning of “dirt”, one must not only know what “material” properties dirt typically has, but also something about the purpose of the word. It is through grasping the purpose of the word “dirt” that the competent speaker understands why custard on a plate is not dirt, while custard on a shirt is. As I will explain in more detail later in this Section, this knowledge allows the competent speaker to distinguish between correct and incorrect applications of the word “dirt”.

For Kovesi, these two different ways of *misunderstanding* the meaning of a word highlight that there are two distinct aspects to *understanding* the meaning of a word. His view is hence incompatible with any account on which only *one* of these aspects is involved with understanding the meaning of a word. Kovesi believes that many meta-ethicists attempt to understand moral language with reference to “material elements” only, at the expense of accounting for the formal element, i.e. the purposes of the words and language, and

considers this to be a significant error (see e.g. Kovesi 1967: 24). In the remainder of this Chapter, I will argue that Kovesi is right on this count.

Indeed, my argument will be that the Dichotomous Claim makes sense only when one is thinking about the meaning of descriptive words *exclusively* in terms of the “material elements”. When we begin thinking about the meaning of such words both in terms of their recognitors *and* their purposes – as we should, if Kovesi’s “purpose thesis” is correct - the Dichotomous Claim will no longer make sense.

I want to explore, then, whether there is genuine tension between Kovesi’s “purpose thesis” and the Fact/Value Dichotomy, and what the implications of this might be for that Dichotomy. Kovesi himself seems aware of the possible tension here, writing that

Indeed, decisions and attitudes, insights, wants, needs, aspirations and standards do enter into our moral notions. But they do not enter from the top, they are part and parcel of our notions [quite generally] ... evaluation is not an icing on a cake of hard facts ... standards, needs and wants also enter into the formation of terms that we usually call descriptive terms (1967: 25).

While Kovesi’s statement that “evaluation is not an icing on a cake of hard facts” is unfortunately opaque, we get the sense from this passage that Kovesi opposes the Johnston-style idea that descriptive and evaluative language are entirely separate from one another. One difficulty here is that since Kovesi does not distinguish between the Humean and Dichotomous Claims, he does not clearly state which form of the Fact/Value Dichotomy he opposes, if any.

However, if we try to put the jigsaw puzzle pieces together here, we can understand the dialectic as follows. Kovesi is asserting that “standards, needs and wants”, etc. enter into *both* our descriptive and evaluative language. That is, for Kovesi, both descriptive and evaluative language must be understood through its “formal element”, purposes, and all of the associated factors relating, for instance, to human “standards, needs and wants”. In

denying that evaluation is “an icing on a cake of hard facts”, and stressing that the “formal element” is also crucial to our descriptive language, Kovesi is opposing a rival philosophical view on which “standards, needs and wants” do *not* enter into the formation of our descriptive language. On this rival view, descriptive language simply gives “hard facts”, separate from evaluation. If we sided with this rival view, we may well end up thinking that descriptive language and evaluative language are fundamentally different in kind, since we will then have the impression that evaluative language must be understood through its “formal elements”, i.e. purposes, the relevant standards and needs, etc., while descriptive language need not be. Kovesi opposes this view, arguing that *both* descriptive and evaluative language must be understood through its “formal elements”, and so on. Kovesi’s view there is no reason to think there is a difference in kind here.

So, if the rationale for thinking that descriptive and evaluative language are different in kind is rooted in the idea that descriptive language need not be understood through a “formal element”, Kovesi will reject that rationale, since he thinks that both descriptive and evaluative language must be understood through its formal element. So while I am not at present arguing against the Fact/Value Dichotomy, I have detected a potential source of tension between Kovesi’s view and that Dichotomy, and I suggest this warrants further investigation.

Kovesi’s “purpose thesis”, however, requires careful handling, if we are to keep to the strategy I outlined in the first Section of this Chapter. One way of using it would be to argue that all language is *essentially* such that it must be understood through its purpose, and on that basis to claim that all language is essentially semantically “thick”. This would be to use the “purpose thesis” to provide a Kovesian version of Lovibond’s approach, which I analysed in Chapter 2, Section 3. Kovesi may have committed himself to this essentialist approach by writing that the “formal element” is “part and parcel” of our concepts (1967: 25), but I think we should not follow him in this regard, since I argued in Chapter 2, Section 3 and Section 1 of this Chapter that it would be advantageous to adopt an alternative, non-essentialist approach.

To apply Kovesi's "purpose thesis" without committing to essentialism, I shall not argue that the "purpose thesis" applies to all possible language, or claim that all language is essentially such that it must be understood through its purpose. As an alternative, I consider Kovesi's "purpose thesis" to provide a particular way of understanding the meaning of language. That is, on Kovesi's approach, the meaning of a given piece of language is to be understood both through what he calls its "formal elements" and its "material elements". If this approach to understanding meaning does not allow us to make sense of the notion of semantically thin descriptive language, then the advocate of the Dichotomous Claim will be left needing an alternative approach with which to make sense of the notion of semantically thin descriptive language. The viability of the Dichotomous Claim will then depend on whether such an alternative approach is forthcoming. If one is not, there would still be nothing to rule out the possibility that the notion might be successfully made sense of in future. Equally, however, there is no reason to accept that semantically thin descriptive language is possible until the notion of semantically thin descriptive language has been explained and made sense of, just as there is no reason to accept that a logically private language is possible until the notion of logically private language has been made sense of (see Chapter 3, Section 1, above).

I also want to note that, while I have explained and sometimes applied Kovesi's terms "material element" and "formal element", that distinction will not be crucial in the argument below. While it was useful to discuss this distinction for the sake of clarifying how understanding the purpose of a word can be helpful or even crucial for understanding its meaning, I am hence not committed to Kovesi's distinction between "material elements" and "formal elements". It is, rather, the "purpose thesis" which plays an important role in the argument below.

I will now turn to applying the purpose thesis to the Humean and Dichotomous Claims, in turn.

Part iii: The Implications of the Purpose Thesis for the Fact/Value Dichotomy

In this part of the present Section, I am going to argue that the purpose thesis undermines both the strong form of the Fact/Value Dichotomy, the Humean Claim, and the weaker form of that Dichotomy, the Dichotomous Claim. I will start by answering the Humean Argument in favour of the Humean Claim (an argument which I first introduced in Chapter 1, Section 2, part iii). Then I move on to my argument against the Fact/Value Dichotomy.

We can use the ideas from Kovesi (1967) introduced above to undermine the *Humean Argument*, which Johnston (1989) uses to motivate the Humean Claim. Johnston tries to motivate the idea that there is a “fundamental logical difference” between descriptive and evaluative language by contrasting paradigmatically evaluative words like “right” or “wrong” with supposedly paradigmatically descriptive words such as “red”. For Johnston, the meaning of a descriptive claim like “there is a red book on the desk” is to be understood in terms of the picture or state of affairs which it represents. But while it seems quite easy to picture the state of affairs which the descriptive “there is a red book on the desk” is supposed to represent, it is very hard to picture any state of affairs which the evaluative “eating meat is wrong” might represent. The result is an apparent crucial difference in how the meanings of descriptive and evaluative claims are to be understood, and this gives the impression that descriptive and evaluative claims have two different kinds of meaning, thus supporting Johnston’s Humean Claim.

Kovesi undermines the Humean Argument by pointing out that colour-words like “red” are *not* representative of all descriptive words (1967: 9). If all descriptive words must function like “red”, Kovesi suggests, it would be hard to say “table” is a descriptive word either. This would mean we could not say a sentence like “this is a table” is a descriptive sentence, which would be very counterintuitive.

The reason for this is that, as Kovesi argues, “table” does not function in the same way as “red”. In the case of “red”, it is easy to picture what must be the case for a simple

descriptive claim like “this patch is red” to be true, since whether something counts as red depends only on the presence or absence of the single observable property we call “red” (Kovesi 1967: 6). With the word “table”, by contrast, it is much harder to say which properties an object must possess if it is to count as a table, with the natural result that it is much harder to picture what must be the case for a simple descriptive claim like “this is a table” to be true. For that reason, it is comparatively easier to argue that words like “table” and “good” must be understood with reference to the purpose of the word, as Kovesi suggests, than it is to argue that “red” must be understood with reference to its purpose.

Of course, we can, as we saw in part ii) of this Section, list various properties that are *typical* of tables: they usually have legs, flat surfaces, and are of roughly waist height, for example. We can, moreover, grant that awareness that these are typical properties of tables is an important part of what it is to understand the word “table” (see part ii) of this Section, above). And therefore it might even help us explain the meaning of “table” (to a child, for instance) to draw a picture of a typical-looking table. But none of these typical properties (flat surface, legs, etc.) of tables are *definitive* of what it is to be a table in the same way it is definitive of a red thing that it has the property we call “red” (Kovesi 1967: 6-7). Tables typically have flat surfaces, but need not, just as they need not have legs (we can imagine a table that is suspended from the ceiling, for instance). This is why it is harder to capture the meaning of “table” *purely* through giving some picture or state of affairs which supposedly corresponds with it than with “red”.

The implication of this difference between colour-words like “red” and other descriptive words, like “table”, is that typical evaluative words, like “good” and “wrong”, are less different from descriptive words than they might initially appear. For in the case of the descriptive claim “this is a table”, the best we can do is imagine a picture which features an object which has the *typical* properties of a table. But then, as I will expand upon further in Section 3, below, there are equally typical properties of “good” and “wrong” things and actions, with which we could generate a “picture” to represent a claim like “this action is good”. For instance, “good” actions typically have the property of increasing the amount of happiness in the world, and the property of being intuitively-approved of by persons (see Meynell 1971: 118-119, I discuss Meynell’s views further in Section 3 of this Chapter).

As Kovesi comments, the contrast between colour-words and evaluative words like “good” often gives philosophers (such as Johnston 1989) the impression of a fundamental difference between descriptive words and evaluative words, but ultimately the same contrast can be found between colour-words and words like “table”, which seem descriptive rather than evaluative. It is therefore a mistake to think that this contrast between colour-words and evaluative words establishes that there is a fundamental semantic difference between descriptive and evaluative language.

The Humean Argument put forward by Johnston, then, appears to rest on a “one-sided diet” (to use a phrase from Wittgenstein 1953: Section 593) of examples of descriptive words. If we focus only on descriptive words like “red”, it looks as though there is a difference in kind between “red” and evaluative words like “good”; a difference which it is tempting to say lies at the semantic level (see Johnston 1989 and Chapter 1, Section 2, part iii), and Chapter 2, Section 1, above). But as Kovesi demonstrates, words like “red” are not representative of the class of descriptive words. If the defender of the Humean Argument is to be consistent, she should say that the contrast is ultimately between words like “red” and other words including both seemingly-descriptive words like “table”, “desk”, etc., and seemingly-evaluative words like “good”, “wrong”, and “duty”.

Depending on whether the advocate of the Humean Argument accepts the Humean Claim or the Dichotomous Claim, they could respond to this argument in different ways. As with previous objections considered in Chapter 2, Section 1, the Dichotomous Claim seems to fare better than the Humean Claim at handling this problem.

The advocate of the Humean Claim could respond by insisting that “table” is not, after all, a descriptive word, but words such as “yellow” and “square”, which are more easily “picturable”, are descriptive. This response, however, would not be very plausible.¹⁰⁴ What

¹⁰⁴ Below, I also consider the slightly different response, according to which “table” is semantically thick while “yellow” is not. The Humean, however, cannot appeal to the notion of semantically thick language, since advocates of the Humean Claim by definition cannot accept the possibility of semantically thick language (see Chapter 2, Section 1, above).

is the Humean to say about sentences such as “there is a table over there”? According to the Humean Claim, all claims which are either descriptive or evaluative must be one or the other, and cannot be both. A Humean who denies that “table” is a descriptive word would, it seems, have to say that “there is a table over there” is evaluative, not descriptive. But to reject that this sentence is descriptive seems very implausible (even if there is a case for saying it is both descriptive and evaluative), so the Humean has no good reply to the present objection.

By contrast, the advocate of the Dichotomous Claim can respond in a more plausible way. I now want to explore this response to set up my argument against the Dichotomous Claim, which I present in the next part of this Section.

The contrast we found between words like “red”, on the one hand, and words like “table” and “good”, on the other, should be explained (according to this response) by saying that “red” and other similar words are “thin” descriptive words, whereas “table”, “good” and other similar words are either thick or “thinly” evaluative. This fits with Parfit’s (2011) claim that shape-words and colour-words are examples of “thin” descriptive words (see Chapter 2, Section 1, above). While it may seem a little counterintuitive to count words such as “table” as thick, we found above that the Dichotomous Claim is flexible about how much of our ordinary language is thick (see Chapter 2, Section 1). So this response seems to be a potentially viable way of defending the Dichotomous Claim.

Let us pause at this point to take stock. From Kovesi’s point of view, it seems, *both* descriptive and evaluative language must be understood through both their “material elements” and “formal elements”. For that reason, Kovesi thinks that descriptive and evaluative claims share the same semantic form, and he is hence suspicious of any view which claims that they are semantically different in kind. To support his view, Kovesi undermines the Humean Argument by showing that the same contrast between words like “red” and words like “good” also appears between words like “red” and other descriptive words like “table”.

While this causes serious trouble for the Humean Claim, the advocate of the Dichotomous Claim seems to have a good response. Recall that according to Kovesi's "purpose thesis", to understand the meaning of a word, the speaker must have due awareness of the *purpose* of that word (cf. Dummett 1973: 295-298, Queloz 2019: 1124-1126).¹⁰⁵ For Kovesi, this semantic feature is "part and parcel" of our notions quite generally (1967: 25). We might wonder, however, whether this is true. An advocate of the Dichotomous Claim, like Parfit, may argue that the "purpose thesis" is accurate for only *some* of our words. If there are some possible or actual words which the purpose thesis does not apply to, this would give a reason for thinking that such words are semantically different in kind from other words, for they (unlike others) are not to be understood through their purpose, or "formal element".

The advocate of the Dichotomous Claim would, therefore, answer Kovesi by claiming that only *thick* words and evaluative words conform to Kovesi's model. This response concedes that Humeans like Johnston are wrong to ignore the salient differences between "yellow" and "table", and also wrong to think that the meaning of "table" can be captured in the same way as the meaning of "yellow". But while, perhaps, Kovesi's discussion of formal elements reveals that *some* descriptive words need to be understood through their purpose, will that be the case for *all* concepts?

If some descriptive words can be completely understood independently of their purpose, Kovesi's purpose thesis is false. Moreover, it seems plausible that apparently-thin descriptive words, like "yellow", are not captured by the "purpose thesis". In fact, if the purpose thesis applies only to thick words, this might yield a good way of *explaining* the logical difference between "thin" descriptive words and other words, like "thick" words.

This helps answer a crucial question for the likes of Parfit (2011): *what features of supposedly "thin" descriptive words, if any, make them logically different in kind from*

¹⁰⁵ Dummett motivates this idea with the following analogy. Just as, to master chess, one must know not only how the pieces are allowed to move, but also that the *point* of the game is to checkmate the opponent's king, so it goes for mastery of concepts. Dummett's idea is that "describing the usage of a concept is not enough to master it; one has to grasp the *point* of the concept", as Queloz explains (2019: 1126).

thick words and thin evaluative words? Why does Parfit think that words such as “yellow”, “square”, “tall” etc., fall into a distinct semantic category? In other words, *do* we have good reason to think that “yellow” is different in kind from “table”?

While Parfit does not address these questions, I believe there is a clear answer. The apparent special feature of words like “yellow” and “square” is that it seems that all one needs, to understand those words, is to be able to recognise, for example, squares (or yellow things) and distinguish them reliably¹⁰⁶ from non-squares (or non-yellow things), through awareness of the appropriate “recognitors”. Consider the example of the school pupil whose teacher draws a square on the whiteboard and says “this is a square!”. Through the teacher repeatedly pointing out to the pupil which shapes are squares and which are not, we can imagine that the pupil will eventually become able to correctly point out squares and avoid pointing out non-squares when given the appropriate prompt. Unlike the case of “table”, in the case of “square” there is a simple test for whether a given object is square: if it has four sides of equal length, it is square. There is no equivalent simple rule for determining what is a table and what is not, which makes it more plausible that understanding of the *purpose* of the word “table” is required to know the meaning of that word. But once the pupil can reliably discriminate between squares and non-squares in the way described, should we not conclude that they have learned the meaning of the word “square”?

If that *is* all it takes to know the meaning of “square”, we will have a counterexample to Kovesi’s purpose thesis. For this pupil, from what we have said about the hypothetical scenario, knows nothing about the *purpose* of the word “square”. All they have is this ability to reliably point out square shapes, when prompted to point to the “square”. The idea is that which Wittgenstein summed up (but did not endorse): “Once you know *what* the word signifies, you understand it, you know its whole application” (1953: Section 264).

¹⁰⁶ I say “reliably” because we clearly do not need to be infallible at distinguishing squares from non-squares in order to have mastered the concept “square” and to know the meaning of the word. Even totally competent users of the word may sometimes mistake a rectangle for a square, for instance (if the rectangle looks similar to a square on account of its sides all being of similar length).

Specifically, however, I am now considering a restricted version of that view, on which this is true only for “thin” descriptive words.

We can imagine the pupil being taught the meaning of “yellow” and “tall” in similar ways to how they are taught “square” in the example above. First, the pupil could be trained to point to yellow-coloured things when prompted, and to avoid pointing to non-yellow things. Second, the pupil could be trained to point out the tallest object in some group of objects, when given the appropriate prompt. The common theme in these cases is the apparent possibility of teaching the meaning of the word by training the pupil to discriminate between things that we *do* use the word to refer to, and things we do *not* use the word to refer to. And this understanding of what it is to know the meaning of a word looks much more plausible in the case of words like “square”, “yellow” and (perhaps to a lesser extent) “tall”, than it does in the case of words like “table”.

So, the likes of Parfit, who wish to defend the Dichotomous Claim, may use this apparent disanalogy between words like “square” and words like “table” to support the view that words of the former kind, but not the latter, are semantically “thin”.

Moreover, on this approach, thin descriptive words *would* end up being logically different in kind to thin evaluative words – vindicating the Dichotomous Claim. For it is difficult to imagine how the meaning of thin evaluative words like “good”, “right”, “duty”, etc. could possibly be taught, or understood, in the same way as “square” is apparently understood by the pupil in our example. How could the meaning of the word “good” be taught through ostension, as the meaning of “square” apparently can be? The advocate of the Dichotomous Claim concludes that it cannot be, since it is of a different logical type. On this view, then, Kovesi’s “purpose thesis” applies only to thick descriptive words and evaluative words, not thin descriptive words.

I have asked what so much as *gives the impression* that so-called “thin” descriptive words are logically different in kind from so-called “thick” words. Having not found answers to

this question elsewhere, I have tried to answer it myself. The idea appears to be that thin descriptive words, unlike thick words, can be fully understood without awareness of the purpose of the word. The apparent motivation for thinking this way is that the words on Parfit's list of "thin" descriptive words seem to be well suited to the analysis Wittgenstein summarises: "once you know *what* the word signifies, you understand it, you know its whole application". On the face of it, this idea seems rather plausible: it *does* seem to pinpoint a genuine semantic difference in kind between thin descriptive words, on the one hand, and thick words and thin evaluative words, on the other. But does this response stand up to critical scrutiny? I will now argue that it does not.

We have seen that the proponent of the Dichotomous Claim needs some account to explain *why* thin descriptive words are logically different in kind from thin evaluative words. If the purpose thesis does not apply to thin descriptive words, this would form the basis of such an account. Kovesi has arguments directed against this strategy. He argues that denying the purpose thesis, even for words like "yellow" and "square", is untenable. As I will now argue, Kovesi is right about this. This leaves the proponent of the Dichotomous Claim lacking a rationale for their claim that thin descriptive words are logically different in kind from thick words and thin evaluative words.

Part iv: My Argument Against the Dichotomous Claim

Kovesi gives the example of a teacher drawing a particular "pointed projection" on the board in front of a class of pupils and saying "this is a tak!" (Kovesi 1967: 39). As with the "square" example above, the teacher trains these pupils to recognise taks and to be able to reliably distinguish "tak" shapes from non-"tak" shapes. But would the pupils thereby have learned the meaning of a new word, "tak"? Kovesi admits that the case can give the impression that this is so. But the impression is, he concludes, misleading.

Here is how Kovesi explains why he rejects the idea that the pupils have learned the meaning of a new word, "tak":

There is... something strange in assuming that the pupils can leave the classroom and say: "Now we know the meaning of the word "tak"... if they see [a tak], what should they do then? Should they stop and say "tak" each time, or point it out to someone saying: "Look here is another tak", or should they perhaps make a record of the number of taks they see each week? What is the point of the word "tak"? The instructor gave no suggestion about this, and consequently did not teach the pupils the meaning of the word. (1967: 39-40).

The point is that it is entirely mysterious, as the case has so far been described, *when* and *why* the pupils, or any would-be user of the word "tak", should use the term. Accordingly, Kovesi claims, the word lacks a meaning – or, at least, its meaning has not yet been explained (1967: 40). We cannot simply assume that it would be correct for the pupils to say "tak" whenever they perceive a tak-shape. It is not as though we only speak of squares when we happen to see one, nor do we exclaim "square!" every time we see a square (Kovesi 1967: 40). So why should it be any different with the word "tak"?

In our hypothetical situation, although the pupils have been taught to associate the word "tak" with a particular pointed shape, we have no information about the purpose or function the word is supposed to serve, or what role it is supposed to play in their lives. Why is Kovesi so adamant that, without any inkling as to the purpose of the word, the pupils have not been taught the meaning of a new word, "tak"?

The answer, I argue, lies in problems which arise from there being, in the case of "tak", no clear answers to the questions of *when and why* the term *should* be used. My suggestion is that, in the absence of any agreement between speakers about when and why a word should be used, we will have no way of distinguishing between a correct and a mistaken application of a word. But if we have no way of distinguishing between a correct and a mistaken application of a word, how can we distinguish between persons who have become *competent* in the use of that word, and persons who have not?¹⁰⁷ For surely the

¹⁰⁷ Cf. Wittgenstein (1953), Section 258. For commentary on and explanation of Wittgenstein on this point see e.g. Schoenbaumsfeld (2016), Ch. 3.

only way to distinguish between those who understand a word, and know its meaning, and those who do not is to draw a distinction between those who are reliably competent at applying the word in the correct ways, and those who systematically fail to so apply the word (such as infants who are just starting out with language)? Without any notion of when and why “tak” should be used, then, we will be unable to distinguish between those who genuinely understand the word and those who do not. In such a situation, it would be empty to speak of the word “tak” having a meaning, or of the pupils in our classroom having learned the meaning of “tak”.

To use a phrase from Wittgenstein’s discussion of “private” language, we have no “criterion of correctness” for the application of the word “tak” (Wittgenstein 1953: Section 258). But without such a criterion, how can we meaningfully speak of a “correct” use of the word in question? Without being able to meaningfully speak of a correct or incorrect use, what basis is there for saying that the pupils have learned the meaning of a new word, “tak”? For it only makes sense to say that the pupils understand the meaning if we have some notion of what it is to use the word “tak” correctly and incorrectly!

None of this is to deny that the pupils might *find* a use for the word “tak” (cf. Kovesi 1967: 40). For example, imagine they go out into the world and find that the presence of “tak”-shaped objects correlates with the presence of a dangerous animal. In this case, they could develop a convention of always shouting “tak!” when they see that shape, with the purpose of warning other people about the risk of danger. If this purpose was generally agreed upon, the pupils could begin to distinguish correct from incorrect uses of “tak”. There would now be a reason for the pupils to, in general, say “tak” when they saw a “tak” shape, and they would generally consider it a mistake to say “tak!” when no tak-shapes were visible to them (since this would be a false alarm). So introducing a purpose for the word grants the possibility of distinguishing correct from incorrect uses, and hence a meaning. This is why I believe Kovesi’s purpose thesis applies even to words like “square”, which can *seem* semantically “thin”, but ultimately are not.

It is not that associating the word “tak” with a particular shape, as the teacher does, is a pointless exercise with no legitimate role in language-learning. It is rather that we tend to forget that this sort of association is not enough *on its own* to give meaning to a word. For a word to have a meaning, according to this approach, it is not enough that speakers associate it with a certain shape, colour, etc.: it must also have a purpose.

How does this apply to the case of the word “square”? The point here would be that it is not as though the teacher’s mechanism of pointing at a square and saying “this is a square!” is enough *on its own* to ensure the pupils learn the meaning of “square”. As Kovesi highlights, competent users of the word “square” do not simply say “square” whenever they see a square. The word “square” rather has a particular use within our language and activities, relating to its purpose in fields like mathematics and art. For example, when children are taught about squares, they are, after all, being taught with particular purposes in mind. When a teacher ostensibly explains the word “square”, she is teaching the pupils a word which *does* have specific purposes, many of which she is aware of and form part of the background as to why the teacher chooses to teach them this word. While she may not explicitly mention those purposes to her pupils, some understanding of the purpose of the word appears to be required if a pupil is to count as knowing the meaning of the word, since it is only with that that the child will have a sense of when and why the word is called for.

While the ability to reliably distinguish squares from non-squares is an important part of becoming competent in the use of the word, then, it is difficult to see how this could be enough on its own. Part of learning the meaning of the word must involve learning the difference between correct and incorrect uses of the word, and this goes beyond merely learning to pick out the appropriate shapes, colours, etc. To understand when it is correct and incorrect to use a word, we need to understand the word’s purpose, which gives us something against which to judge whether an application was correct or incorrect.¹⁰⁸ A pupil who can reliably distinguish between squares and non-squares, for example by pointing, may be unable to use the word “square” appropriately in a conversation. In such a

¹⁰⁸ There are affinities between the Kovesi-inspired view I express here and McDowell’s (2000: 40-43) criticism of the “mechanistic” view of understanding of linguistic rules; cf. Ginsborg 2018).

case, we would say the pupil does not know the meaning of the word. What is missing here is an appreciation of the purpose of the word, which helps competent speakers and language-learners to start applying the word correctly.

The problem for the Dichotomous Claim is that, let us recall, it requires an account to explain and justify its claim that thin descriptive words are logically different in kind from thin evaluative words. Based on Parfit's list of thin descriptive words, which includes "red", "square", "tall" etc., I supposed that the proponent of the Dichotomous Claim would try to argue that the distinctive logical feature of such words is that they can be fully understood *without* awareness of their purpose. For in the case of such words, it is easy to get the impression that all we need, to understand them, is to be able to reliably recognise the things that they refer to (e.g. squares, colours, tall things, etc.).

Unfortunately for the Dichotomous Claim, however, we have found that impression to be misleading. Upon closer analysis, it seems that even for such apparently "thin" words, an understanding of their purpose is crucial for understanding their meaning. Without the word having a purpose, we do not have a way of distinguishing between a competent and incompetent user of the word, I have argued. And without that, it is difficult to make sense of the idea that the word has a meaning, as Kovesi's "tak" example illustrates. Since the notion of semantically "thin" descriptive language appears to rest on the idea that the meaning of a word can be understood without reference to its purpose, we hence have good reason to be sceptical about whether the notion of semantically "thin" descriptive language is intelligible. What I suggest is that the idea that semantically "thin" descriptive language is possible rests on a conceptual confusion: that words can have meaning even if they lack a purpose.

One objection that could be made against my argument runs as follows. Is it not the case that the purpose of concepts such as "square" is to give descriptions? If this is so, then in a sense it is still a "purely descriptive" concept, because its role in our lives is, after all, to describe things.

My first response is to say that the Dichotomous Claim is, as I explained in Chapter 2, Section 1, a *semantic* thesis on which it is possible to have thin descriptive language which is different in semantic kind from thin evaluative language. Pointing out that the function of certain words is to describe things is not the same as saying that those words have a special kind of meaning. Rather, since they must be understood through both their “recognitors” and their purpose(s), supposedly “thin” descriptive concepts still seem to be, semantically, of the same kind as concepts like “table”, or even non-controversially thick concepts like “brave”, “lewd”, etc. My point is that, even *if* the purpose of “square” and similar concepts is “purely descriptive”, my argument still undermines the attempt to make sense of the idea of “thin” descriptive language under consideration here, because it will remain true that concepts like “square” cannot be fully understood without reference to their purpose.

For the objection to succeed, one would have to show not only that the purpose of words like “square” is purely to describe, but also that this fact implies that the concept is semantically “thin”, and I am sceptical about both these points, for the reasons just outlined.

My second response to the objection is to say that, even *if* a certain word is exclusively used for describing things, the question we need to ask is what *makes* this the case? The most plausible answer, I suggest, is that in such a case, the speakers’ *purpose* is for the word to describe things. But this means that, after all, the word still has a purpose, and competent users of that word will be competent at applying the word in part because they understand its descriptive purpose! So what might appear to be a counterexample to the Kovesian approach to understanding the meaning of words is in fact, on the contrary, neatly captured by that approach.

The argument I have just given explains why the Dichotomous Claim is appealing *prima facie*. It explains why it looks as though some descriptive words are thinly descriptive, like “square” and “yellow”. These words appear to be logically different in kind from thick words and thin evaluative words like “good” and “duty” because we have the impression

that being able to distinguish, say, squares from non-squares, is enough to understand the word “square”, but that the same model cannot apply to words like “good” (or indeed “brave”). The fact that words like “square” are often taught through ostension makes this impression even more enticing, as Kovesi notices (1967: 39). It is not surprising, then, that many philosophers endorse the Dichotomous Claim. This is my explanation of why the Dichotomous Claim is so widespread. But as I have just argued, this picture rests on a mistaken view about what it takes for a word to be meaningful.

Part v: What Does My Argument Achieve?

I will now relate the argument I have just given back to the strategy I set out in Section 1 of this Chapter. As I wrote in that Section, I am not aiming to provide a final refutation of the Fact/Value Dichotomy. Instead, I am challenging the advocate of the Fact/Value Dichotomy to *make sense* of the idea of semantically “thin” descriptive language.

I started by considering the claim from Parfit (2011) that semantically thin descriptive and evaluative language is just a mundane reality, not an abstract idea that stands in need of explanation or justification of any kind. For Parfit, words like “red”, “square”, “duty”, and “right” are semantically thin concepts which can be used to construct semantically thin claims. While Parfit’s view looks plausible *prima facie*, I argued that to convince us his view is correct, Parfit would need to pin down what it is that makes these supposedly “thin” concepts semantically different in kind from other concepts.

The apparently distinctive feature of such concepts, to which I suggested Parfit could appeal to support his position, is that it seems that understanding such concepts requires merely that one is able to reliably recognise what the word refers to (e.g. the right kind of object, or the right colour, etc.) This apparent special feature of such terms seems to differentiate them from others, which is presumably why Parfit chose them for his list. I have argued (following from and developing Kovesi’s thoughts), however, that this

apparent special feature is merely apparent. Being able to reliably recognise what a word refers to is, I argued, insufficient on its own to understand the meaning of these words.

Of course, Parfit, or other advocates of the Dichotomous Claim, may propose an alternative way of explaining what they mean by semantically thin descriptive language. If they do so, we would need to judge that fresh attempt by its own merits. My point is simply that until such a further attempt is essayed, we have to make do with considering the only explanation that I can find on their behalf, and I do not think that explanation succeeds.

My argument does not only cause trouble for Parfit's claim that many words in ordinary language are semantically thin descriptive words. It also causes trouble for the claim that it is possible for there to be a semantically "thin" descriptive language.¹⁰⁹ For, if the advocate of the Fact/Value Dichotomy is to give any substance to this idea of semantically "thin" descriptive language, they need to either give us examples of it, or explain to us what special features that language would have which would differentiate it from other kinds of language. As things stand, I do not see what these special features could be. The main candidate that I can see has now been found to be inadequate, since it rests on the mistaken view that one can understand supposedly "thin" descriptive concepts without an understanding of the purpose of those concepts.

My argument therefore indicates that, contrary to appearances, there is no genuine semantic difference in kind between "red", or other supposedly "thin" descriptive words, and "right", or other supposedly "thin" evaluative words. Parfit and other defenders of the Fact/Value Dichotomy are hence left in need of some alternative way of making sense of (i.e. making substantial) the idea that semantically "thin" descriptive language is possible. It remains entirely open to them to find some such way of substantiating their idea, and, of course, if they find a cogent response to my questioning, then I shall readily give up the

¹⁰⁹ As I established in Chapter 3, Section 1, this feature of my argument makes it significantly different and, I believe, preferable to the arguments of Putnam (2002) and Lovibond (1983) (see my account of the disadvantages of their approaches in Chapter 2, Sections 2-3).

cause. I confess, however, I find it hard to imagine how they could substantiate their idea, given what we have discussed about meaning.

For evaluative concepts, including those which are widely considered to be semantically “thick”, it is widely appreciated that an understanding of their purpose is necessary if one is to (fully) understand their meaning. The surprising result from Kovesi’s (1967) discussion is that this turns out to be true, also, of concepts like “table” and even, less obviously, concepts like “square” and “yellow”. Hence, I believe, our default view should be to deny that these words are of different semantic kinds to each other.

My argument satisfies the three desiderata laid out in Chapter 2, Section 3. First, it challenges both the Humean Claim and the Dichotomous Claim and hence meets the *Completeness* desideratum. Second, as I will now explain, it does not depend on any kind of essentialism.

The reader may think that, since my argument relies on the purpose thesis, it does after all depend on a kind of essentialism. *Does* all possible meaningful language necessarily conform to the purpose thesis? Rather than answering “yes”, I have instead argued that the notion of language that does not conform to the purpose thesis has not yet been satisfactorily explained: its grammar, to borrow a phrase from Wittgenstein (1953), has not yet been explained to us. To refer back to my comparison with Mulhall’s interpretation of the private language argument in Chapter 3, Section 1, I allow it is possible that such an explanation could, in future, be offered, and could vindicate the Fact/Value Dichotomy. Hence, I allow that there *could* be exceptions to the purpose thesis, in principle, which means I do not commit to any form of essentialism. However, my argument still cuts against the Fact/Value Dichotomy, at least *until* or *unless* some vindicatory explanation of the idea of semantically “thin” descriptive language is provided. I therefore recommend that we *provisionally* reject the Fact/Value Dichotomy, unless or until a cogent vindicatory explanation is offered which gives its advocates a way of responding to the argument I developed above. My argument hence meets the *Non-Essentialism* criterion.

Finally, my argument has pinpointed the assumptions about language and meaning which seem to underpin the Fact/Value Dichotomy. Specifically, I have suggested that it is not “pictorial semantics” (as suggested by Putnam 2002) or the “metaphysically heterogeneous view of language” (Lovibond 1983) which are responsible, especially when we consider Parfit’s (2011) defence of the Dichotomous Claim. Rather than being based on such theories, the Dichotomous Claim, I argued, is based on the seemingly innocuous idea that all it takes to know the meaning of some basic words like “square” and “yellow” is to be able to reliably identify which things are square, or yellow, etc. Although the idea seems innocuous, however, I argued, drawing from Kovesi (1967), that this view is based on a serious misunderstanding of the meanings of such words, and what it takes to understand such words.

Conclusion to Chapter 3, Section 2

By drawing from Kovesi, I have undermined the Fact/Value Dichotomy. The Dichotomous Claim rests upon the idea that “thin” descriptive and “thin” evaluative words are logically different in kind. My strategy has been to challenge the Dichotomous Claim by pressing for an explanation of the meaning of purportedly “thin” descriptive words which preserves their thinness. Can we make sense of the idea that these words have a meaning which is different in kind from that of evaluative words? On the face of it, this seems possible, since it seems that to understand the meaning of “thin” descriptive words, all one needs is to be able to recognise what the word refers to (e.g. squares, the colour red, etc.) As Kovesi argues, though, this is not sufficient, on its own, to make such terms have a meaning. This leaves the proponent of the Dichotomous Claim lacking a way of making sense of their idea of thin descriptive language which is different in semantic kind from thick and other evaluative language. Until or unless such an explanation is provided, I argue we should reject the Dichotomous Claim.

In this Section, I have rejected the Dichotomous Claim. But aside from my direct challenge to the Dichotomous Claim, there is a further way of persuading the reader that we ought to reject the Dichotomous Claim. If I can show that rejecting the Dichotomous Claim has

beneficial consequences for meta-ethics, this could add further justification for my position. I now turn to why I believe rejecting the Dichotomous Claim has important benefits for meta-ethical theory, especially for Naturalist realism.

Chapter 3, Section 3: Meta-Ethical Naturalism Without the Dichotomous Claim

In part i) of this Section, I will explain why many dominant Cognitivist views in meta-ethics are heavily shaped by, and depend on, the Dichotomous Claim, which I undermined in the previous Section. I then identify and explain a serious problem which meta-ethical Naturalism faces if it embraces the Dichotomous Claim. Parfit (2011) calls this the “Normativity Problem”. In short, the problem is that Naturalism appears to *explain normativity away* rather than explain it, and Parfit, rightly, finds this objectionable. I then set out an alternative version of meta-ethical Naturalism, which does *not* embrace the Dichotomous Claim. This version of Naturalism, I argue, explains the relevant linguistic data and solves the Normativity Problem. Hence, it is to be preferred to the conventional forms of Naturalism which Parfit considers.

My first task is to survey some popular meta-ethical views which depend on the Dichotomous Claim.

Part i: Cognitivist Meta-Ethical Theories and the Dichotomous Claim

Parfit’s (2011) arguments for Non-Naturalist Cognitivism¹¹⁰ explicitly depend upon the Dichotomous Claim. For Parfit, there is a strict semantic difference in kind between evaluative language and descriptive language. As such, he sets out the Non-Naturalist Cognitivism he endorses as the view that “when... normative¹¹¹ claims are true, they state

¹¹⁰ I follow Parfit, and meta-ethical literature generally, by defining “Cognitivism” as the view that moral claims are truth-apt (see also Jackson 1998: 113). Jackson uses the term moral “sentences” instead of “claims”, but I take there to be no significant difference between these terms.

¹¹¹ Parfit’s preferred terminology is to talk of “normative” language, whereas in this thesis I, following Johnston (1989), use the more general term “evaluative”. However, the differences, if any, between “evaluative” and “normative” language are not my focus here, and the arguments I made should apply equally to both.

irreducibly normative facts” (2011: 266). But what is an “irreducibly” normative fact? As we shall see, Parfit’s way of understanding “irreducibly” normative facts reveals how his Non-Naturalist Cognitivism depends on the Dichotomous Claim.

By “irreducibly normative”, Parfit means that these claims cannot be re-stated in non-normative (what he would call “naturalistic” or descriptive) terms. This is clear from the context, since Parfit distinguishes Non-Naturalist Cognitivism from Naturalist versions of Cognitivism, according to which normative claims *can* be re-stated in non-normative language. “According to Analytical Naturalists, normative words have meanings which can be analyzed or defined by using naturalistic words”, Parfit says (2011: 266). For example, an Analytical Naturalist could claim that the meaning of the evaluative claim “telling the truth is good” is equivalent to “telling the truth maximises happiness”. (This is just one possible example of an Analytical Naturalist claim, not something to which all Analytical Naturalists would have to commit). This is an example of a “reduction” of a normative claim. What happens in the reduction is that the reductionist takes what appears to be a normative claim (“telling the truth is good”) and tells us that, in truth, these words really stand for a non-normative descriptive or naturalistic claim (“telling the truth maximises happiness”). Put in slightly different terms, the reductionist takes a claim which appears to ascribe a *normative* property (“goodness”), and says that the claim really ascribes a *descriptive* or *naturalistic* property (“maximises happiness”) (see Klocksiem 2019).

The Analytical Naturalist, then, takes herself to be substituting some language which is “purely normative” (i.e. semantically “thin” normative language) with some language which is “purely descriptive”, and non-normative (i.e. semantically “thin” descriptive language). This brings out why Analytical Naturalism depends on the Dichotomous Claim: it depends on the idea that there can be semantically thin descriptive language. If my argument against the Dichotomous Claim succeeds, it is a mistake to think that we could replace “thin” normative language with “thin” descriptive language, because we cannot *make sense* of the idea of such “thin” descriptive language.

However, Parfit's version of Non-Naturalist Cognitivism equally depends on the Dichotomous Claim. Recall that, for Parfit, Non-Naturalist Cognitivism involves the claim that there are "irreducible" normative truths. As we can now understand, this means that Parfit thinks there are "thin" normative claims which cannot be logically reduced to non-normative descriptive claims. Parfit explains that, on his view, to believe that there are "irreducibly normative" claims means to think that "these concepts and claims cannot be re-stated in non-normative terms" (2011: 295). Those who reject the Dichotomous Claim and follow my argument in Chapter 3, Section 2, however, should be suspicious of what Parfit means by "non-normative terms". If he means "terms with a meaning which is semantically different in kind from the meaning of normative (or evaluative) terms", then he will be committed to the Dichotomous Claim. Since Parfit explicitly endorses the Dichotomous Claim (2011: 265-266), and he is up front about how this Claim shapes his meta-ethical methodology, it is reasonable to conclude that this *is* what he has in mind by "non-normative terms".

To sum up my claims so far in this part: both Analytical Naturalism and Parfit's Non-Naturalist Cognitivism depend on the Dichotomous Claim. The Analytical Naturalist tries to analyse the meaning of purportedly thin normative claims by reducing them to purportedly thin descriptive claims. Hence, if we should reject the idea of thin descriptive language, we should also reject Analytical Naturalism. Parfit's Non-Naturalist Cognitivism seeks to preserve what it takes to be the strict distinction between thin normative language and thin descriptive language, and Parfit's objections to Analytical Naturalism centre precisely around his claim that it is wrong to violate this distinction (Parfit 2011: Ch. 24). Hence, Parfit's view trades just as much in the Dichotomous Claim as does Analytical Naturalism.

It is also important to address Non-Analytical Naturalism. According to Parfit, Non-Analytical Naturalists accept that there are some irreducibly normative claims (2011: 295). However, Non-Analytical Naturalists, unlike Non-Naturalists, maintain that all facts are natural, and that some natural facts are also normative. On this view, a normative fact is a natural fact which can be stated both using irreducibly normative language and thin descriptive naturalistic language (2011: 295). While, according to Non-Analytical

Naturalists, the concept “right” is, semantically, thinly normative, and there are thin normative claims, all true normative facts can also be expressed by thinly descriptive statements of natural facts. On this view, such facts are both normative and natural (Parfit 2011: 324, for examples of Non-Analytical Naturalism see e.g. Sturgeon 1985, Brink 1986, 1989, for discussion see Suikannen 2016).

Plainly enough, Non-Analytical Naturalism of this form also trades in the Dichotomous Claim. First, it appeals to the notion of “irreducibly normative” language, just as Parfit’s Non-Naturalist Cognitivism does, and for the same reasons given above, this reveals a commitment to the Dichotomous Claim. Second, this Non-Analytical Naturalism also trades in the idea of thinly descriptive, naturalistic language, since it says that true normative facts can be captured by thin descriptive naturalistic statements. So, if we reject the Dichotomous Claim, we should also avoid making the Non-Analytical Naturalist’s claims that there is i) “irreducibly normative” language and ii) thin descriptive naturalistic language (both of which, on this view, can capture normative truths).

On Parfit’s way of defining Analytical Naturalism, Non-Analytical Naturalism and Non-Naturalism, then, each of these views accepts the Dichotomous Claim. The result is that, on Parfit’s set-up, it looks as though all major Cognitivist meta-ethical views accept the Dichotomous Claim. My question is: what would a Naturalist meta-ethics which rejected the Dichotomous Claim look like?

Part ii: Parfit’s Objection to Naturalism: The Normativity Problem

This question is especially urgent because of Parfit’s (2011) objections to Naturalism. Parfit’s major concern about Naturalism is that it seems to *eliminate* normativity rather than elucidate it. He calls this the “Normativity Problem” for Naturalism. The easiest way to understand Parfit’s “Normativity Problem” for Naturalism is to consider the Analytical Naturalist view, which attempts to explain the meaning of normative claims by giving a descriptive claim which the normative claim supposedly “reduces” to upon analysis (for

example, the Analytical Naturalist may say that “telling the truth is morally good” means “telling the truth maximises happiness”). Parfit’s concern about this view is that, because the Analytical Naturalist tries to explicate the meaning of thin normative claims by substituting in thin descriptive claims in its place, the Analytical Naturalist appears not to have explained the nature of normative statements, but rather to have *explained them away*.

Parfit argues that, contrary to Naturalists, a thin normative claim and a thin descriptive claim *cannot* “state the same fact” (2011: 326). Parfit uses his *Burning Hotel* case to support this claim: imagine you are standing on a balcony on the second story of a burning hotel, and your only chance of saving yourself is to jump off. Since your life is worth living, Parfit says (2011: 326), **you ought to jump**. According to some Naturalists, the normative fact “you ought to jump” is the same as the fact that “jumping would do most to fulfil your present fully informed desires, or is what, if you deliberated in certain naturalistically describable ways, you would choose to do” (2011: 326). But, Parfit says, these two sentences just do not have the same meaning (2011: 326), and so cannot state the same fact.

Is it not perfectly coherent, Parfit asks, to think that the person on the Burning Hotel balcony *ought* to jump, even if it is *neither* what they would choose to do, nor what they would desire to, in any circumstances? (2011: 327). Since it could be true that you ought to jump at the same time as it could be false that you desire or would choose to jump, Parfit says, the fact that you ought to jump *cannot* be the same as the fact that you desire to jump. And the same problem will apply for the Naturalist view, it seems, no matter which specific version of Naturalism we consider. The point applies equally to both the Naturalist who says that “you ought to X” states the same fact as “doing X would maximise happiness” and to the Naturalist who says “you ought to X” states the same fact as “X is what, in certain conditions, you would desire to do”. So Parfit’s *Burning Hotel* case generalises in such a way that it threatens not just specific forms of Naturalism, but Naturalism in general.

While part of Parfit's concern with Naturalism is hence that it violates an apparently common-sense distinction between normative and natural facts, his worries do not end there. He imagines a Naturalist doctor who is asked by an ethics committee to consider a scenario where she could secretly kill one of her patients to use that patient's organs to save the lives of five other patients, thereby allowing those five other patients to lead very happy futures they would not otherwise have had. In response, the doctor says "I ought to kill this patient, since this act maximises happiness" (2011: 342). The ethics committee are, understandably, horrified, and consider debarring the doctor from any medical post (2011: 342). Yet, as a meta-ethical Naturalist, the doctor explains:

"When I claimed that I ought to kill this patient, I was only stating the fact that this act would maximise happiness. On my view, that is the property to which the concept *ought* refers. I was not claiming that this act would have some *different* property of being what I ought to do. On my view, there is no such different property. The property of maximising happiness is the *same* as the property of being what we ought to do." (Parfit 2011: 342).

Parfit thinks this example should worry Naturalists. While the ethics committee may have been wrong to think that this doctor has an objectionable moral view, Parfit says, the trouble is that she appears not to have "*any* moral view" (2011: 343, emphasis in original). Given Parfit's view on the logical difference between descriptive and normative (or evaluative) language, it is quite understandable that he draws this conclusion. It seems that, after all, the doctor was, in her own mind, simply saying that killing the patient would maximise happiness. For Parfit, this is just a belief about a natural fact, and cannot be a normative belief of any sort. So while the doctor takes her belief to be normative, Parfit takes it to be a mere belief about a natural fact. And after all, on the doctor's own admission, her belief *is* just a belief about a natural fact! The result is that Parfit considers the doctor to have, at best, pseudo-normative beliefs, that may masquerade as normative beliefs, but turn out on closer inspection to be beliefs about natural facts. So Parfit concludes that the doctor's acceptance of meta-ethical Naturalism leads her to lack any normative beliefs. This is why Parfit thinks that Naturalism does not *explain*, but rather *eliminates*, the phenomenon of normativity, to the point that those who wholly embrace it, like our imaginary doctor, end up without any normative beliefs at all.

I think Parfit is right about this, but with a qualification. I think that versions of Naturalism *which accept the Dichotomous Claim* should be rejected, precisely because, as Parfit says, they end up eliminating the phenomenon of normativity rather than elucidating it. In other words, if the Dichotomous Claim is correct, then Naturalists must, in the end, explain away normativity rather than elucidate it, for the reasons Parfit gives. This would mean that Naturalism would be “close to nihilism”, as Parfit claims (2011: 368). The only way to escape Parfit’s trap is to show that Naturalists *need not* seek to explain normativity away by substituting in a thin descriptive language to replace normative language – because the idea of thin descriptive language is a chimera anyway.¹¹² We should adopt a version of Naturalism which escapes Parfit’s trap by avoiding the Dichotomous Claim.

Part iii: Naturalism Without the Dichotomous Claim

The conceptual link between ascriptions of natural properties, such as “this maximises happiness”, and ascriptions of normative and moral properties, such as “this is good”, is not as strong as the versions of Naturalism considered by Parfit take it to be. However, that is not to say that there is no conceptual link between them whatsoever.

To understand this point, consider Moore’s (1903) Open Question Argument against Analytical Naturalism.¹¹³ Moore’s observation is that it makes perfect sense to say, for instance, that “that action was good, even though it did not maximise happiness”, or “that action was good, even though it was not what that person would have, if fully informed about the relevant facts, desired to do”. From the fact that such sentences make sense, Moore infers that goodness cannot be equivalent to maximising happiness, or to being what the agent would desire to do in certain conditions, or indeed to any other natural property.

¹¹² As I argued in Chapter 3, Section 2, above.

¹¹³ For the influence of Moore on Parfit see e.g. Parfit (2011: 346).

Moore took this to show that meta-ethical Naturalism was a non-starter, leading him to conclude that goodness is a simple, indefinable non-natural property. The conventional Naturalist reply to Moore in recent years has been that Moore's Open Question Argument shows only that goodness is not *conceptually* identical to "maximising happiness" (or any other natural property), but does not disprove that "goodness" and "maximising happiness" (for example) refer to the same property – hence the recent popularity of *Non-Analytical* Naturalism (for more detail see Suikkanen 2016).

I believe we should, however, respond to the linguistic data to which Moore draws attention in a different way. Here is Meynell (1971: 118-119):

Moore considers one by one the definitions of good usually offered; like Napoleon, he prudently engages the enemy forces severally, rather than attacking them all together in combination. Had he done the latter, I believe the outcome of his argument would have been different. It makes sense, to be sure, to claim that an action is (say) good, though it is likely to contribute to the happiness of no one; or that another action is good, though it makes no contribution to man's evolutionary progress (however this is conceived) [etc.] ... but suppose I say that an action is good, though it promotes and is intended to promote no one's happiness, though it is intuitively approved by no one, though it makes no contribution to man's evolutionary progress however conceived, and though it is against the revealed will of God ... do I not in effect virtually reduce to nonsense my original statement that the action is good?

The crucial idea here is that if somebody claimed some action to be good, even though it lacks all of the natural properties Meynell mentions, we *could not understand* their claim that the action was good: hence why Meynell says that their original claim would "virtually reduce to nonsense". I will henceforth refer to this as *Meynell's Observation*.

Meynell asks: if the action had no such qualities (e.g. the quality of producing happiness), on what basis could one possibly coherently believe it to be good? Perhaps the speaker could rectify the situation by pointing to some *other* feature of the action which made it good despite its lacking the usual indicators of goodness. If the feature they draw attention

to has an appropriate link with goodness, they might thereby show that they understood the concept “good” well enough all along. But absent such an explanation, their statement would leave us wondering why they would call the action good, and we would be justified in asking whether this person grasps the meaning of the word “good”.

Although Meynell does not explicitly mention Kovesi in his paper, this is reminiscent of the situation we found when considering Kovesi’s example of the word “table”. While various qualities, such as having a flat surface and legs, are typical of tables, they are not essential for something to be a table. (We cannot deduce that some object is a table just because it has a flat surface, and it would be patently false to claim that “X is a table” *means* “X has a flat surface”). Nevertheless, being aware that those qualities are typical of tables is crucial for understanding the meaning, and correct application of, the word “table”. The same situation appears to obtain in the case of the word “good” and the qualities of maximising or producing happiness, fitting with what people would desire in certain ideal circumstances, etc.

Meynell’s Observation can hence be well-explained with the “material elements” and “formal elements” framework from Kovesi (1967). The reason it makes sense (as Moore rightly says it does) to ask whether an act which promotes happiness is good is that promoting happiness is to goodness much as what having a flat surface is to something’s being a table. A competent speaker knows that tables typically have flat surfaces and that good things typically produce happiness, but equally that tables can very well be tables without flat surfaces and good things can very well be good without producing happiness. However, the speaker who is unaware that good things typically produce happiness is not just making a kind of factual mistake, but displaying an imperfect grasp of the concept “good”.

The more typical qualities of goodness one denies an action has, the less sense it makes to continue calling that action good – and at some point, a statement of the kind Meynell considers becomes unintelligible, and casts doubt on whether the person saying it knows the meaning of the word “good”.

Once we understand the relationship between these natural properties (e.g. “producing happiness”, “being intuitively approved by everyone”, etc.) and goodness with Kovesi’s framework, we have already opened up ground between our view and the forms of Analytical and Non-Analytical Naturalism that Parfit considers. Unlike the Analytical Naturalist, we will not say that moral statements like “X is good” can be reduced to statements like “X promotes happiness”, since this would be as wrong-headed as to say that “X is a table” can be reduced to “X has a flat surface”. Moreover, unlike the Non-Analytical Naturalist, we will not claim that “X is good” refers to the same properties, or states the same fact, as statements like “X promotes happiness”, since this would be as wrong-headed as to say that “X is a table” refers to the same properties, or states the same fact, as statements like “X has a flat surface”.

Rather, our picture is now one on which the conceptual connection between “X is good” and statements like “X promotes happiness” is weaker than claimed by the versions of Naturalism Parfit considers. But it is still tangible. As Meynell puts it:

While value-judgments do not logically entail propositions about any one kind of matter of fact, they are not so logically independent of matters of fact that you cannot render a value-judgment so eccentric as to approximate to nonsense by enough factual qualifications (1971: 119).

Meynell’s terminology is slightly different from that I have been operating with, but the essential point remains the same. Moore’s Open Question Argument can give us the impression that there is no logical relation at all between statements of natural facts and moral statements, and many philosophers (such as Johnston) conclude that this is so. But while the logical relation between these kinds of statement is not as close as the conventional Naturalist believes it to be, we have seen that it is present nevertheless. The idea here – building on what Meynell writes - is that **Meynell’s Statement:**

“Action X is good, even though it promotes and is intended to promote no-one’s happiness, is intuitively approved by no-one, does not contribute to the evolutionary progress of mankind however conceived, and does not do the revealed will of God.”

is not an outright contradiction, because the fact that an action possesses those natural properties does not strictly entail that the action is good, and the fact that it lacks those natural properties would not strictly entail that the action is not good. (Compare how Kovesi believes that some object having a flat surface and legs does not strictly entail that it is a table, even though it may indicate, to a competent speaker, that it is a table). Recall also my observation above that the speaker might be able to render the above statement intelligible by adding a further qualification. For example, they might expand:

“While Action X lacks all of those properties, X is what the agent would desire to do, if fully informed about the relevant facts. On my view, this feature is importantly linked to goodness, in the way many people think the feature of promoting happiness is. So I think this property of Action X being what the agent would desire to do, in certain ideal circumstances, suffices to make X good even though it lacks many of the other typical indicators of goodness.”

One may disagree with this claim, arguing that the natural property of being what the agent would desire to do in certain circumstances is insufficient to make an action good. But this would be a disagreement about which natural properties are necessary, or most important, for an action to possess in order to be good. This debate is not my concern here. For our purposes, it is important to see that the person making this claim has, at least, rendered their view intelligible. We may *disagree* with the importance they attach to desires in relation to goodness, but we can now at least understand *why* they claim that action X is good. This is different to the situation before they added this qualification, because at that stage it was mysterious why they were still prepared to call the action good, and hence it seemed probable that their application of the word “good” was based on a misunderstanding.

This reveals that, as Meynell argues, although **Meynell's Statement** is not an outright contradiction, on its own it is unintelligible, and would indicate that the person who uttered it may not fully understand the concept "good". To spell out the logical link between calling something "good" and ascribing to it natural properties such as the kind we have been discussing, Meynell introduces the term "loose entailment":

"Statement α *loosely entails* statements p, q, r and s when to affirm α and at the same time to negate p (or any single one of the others) is not to contradict oneself, though to assert α and at the same time to negate p, q, r and s all together is either to contradict oneself or to talk so eccentrically as to be unintelligible." (1971: 119)

This accurately captures the role played by natural properties in the grammar of moral terms like "good". To recap, my main reason for embracing this view is that it explains the linguistic phenomena to which Moore and Meynell draw attention. First of all, it explains why, as Moore observes, it makes perfect sense to say "X is good, even though it does not promote (or maximise) happiness": this is because competent speakers know that something can be good even if one of the typical indicators of goodness (promoting happiness) is absent, just as we know that something can be a table even if one of the typical indicators of something being a table are absent.

Second, it explains why **Meynell's Statement** is not outright contradictory, but also why it is unintelligible (at least absent some further qualification or explanation). The more relevant natural properties one denies an action or thing has, the less sense it makes to continue calling that action or thing good. There is a spectrum here of degrees of intelligibility that a statement of the form of **Meynell's Statement** can have; as Meynell suggests, such a statement may simply be puzzling or eccentric from the point of view of a competent speaker, or it may be outright unintelligible, and amount to a clear misapplication of "good". Where any given statement lies on this spectrum will be determined by two factors: firstly, how many typical material elements one denies to the thing or action, and secondly, how important the material elements one mentions are in relation to making something, say, good, or a table. For instance, it may be that denying

something promotes happiness makes it, *ceteris paribus*, less intelligible to say that it is good than, say, denying that it is something people intuitively approve of does.

Competency with use of a concept comes in degrees, and so does the lack thereof!

By adopting this view, we hold on to the basic Naturalist insight that there is an important logical connection between ascriptions of moral properties and ascriptions of certain natural properties. At the same time, we take on board Parfit's objections against conventional forms of Analytical and Non-Analytical Naturalism. Since I do not argue that moral claims are reducible to, or state the same facts as, naturalistic descriptive statements, I am not "explaining away" normative or evaluative language in a way that gives rise to Parfit's Normativity Problem. I take this to be a considerable advantage of my view over the other forms of Naturalism set out by Parfit.

My view helps explicate the meaning of moral words by revealing their conceptual link with ascriptions of natural properties. So, for instance, it is plain that things correctly called "good" are typically things which promote happiness. If somebody did not know or understand that good things *generally* promote happiness, we would rightly conclude that they do not properly understand the word "good" (cf. Meynell 1971: 121). This, however, allows us to nevertheless maintain that good things may not *always* promote happiness, and leaves open whether any given good thing has the natural property of promoting happiness.

The reason this view does not "explain normativity away" is that it does not try to substitute normative language with thin descriptive non-normative language, nor does it say that normative claims state the same facts as descriptive, non-normative claims. Rather, my view reveals that descriptive and normative language cannot be completely understood separately from one another.

On my view, to say that some action or thing is good is not merely to positively evaluate it, but moreover it is to say that it possesses some natural property or properties which ground

that positive evaluation. Hence, to call something good is also to signal, to other competent speakers, something about which natural properties we can expect it to have. This is why it is hard to understand a claim like **Meynell's Statement**, which calls something good but then denies that thing has some of the main hallmarks of goodness.

At the same time, to ascribe certain natural properties to some action or thing is also to evaluate it. Consider, as an example, the **Converse Meynell Statement**:

“Action X promotes, and was intended to promote, happiness, it is intuitively approved by almost everyone, it contributes to the evolutionary progress of mankind as ordinarily conceived, and it is in line with the revealed will of God – but Action X is not good”.

Plainly, this would be at least equally paradoxical and difficult to make sense of as **Meynell's Statement**. Once we have described an action in a certain way, such as by ascribing to it the natural properties listed above, it starts to make little sense to deny that it is good. This shows how, in the process of describing the action that way, we have already indicated that we evaluate that action positively, at least to some extent: hence why it would be confusing, and hard to understand, if we follow up this description by saying that the action is not good.

Rather than portraying descriptive and evaluative language as logically separate, then, this form of Naturalism enriches our understanding of both descriptive and evaluative language, by showing how description and evaluation logically tie in with one another.

Let us recall Parfit's example of the doctor standing before the Ethics Committee. Since this doctor accepted a version of Naturalism which endorsed the Dichotomous Claim, when pressed on the matter the doctor admitted that their belief that killing five patients to save one amounted simply to the naturalistic, empirical belief that killing five patients to save one would maximise happiness. As Parfit comments, if the doctor says this, we

should conclude that it is not that the doctor has an objectionable normative belief, but rather no normative belief at all (2011: 343).

Consider what would happen if, instead, the doctor embraced my non-Dichotomous version of Naturalism. When pressed to explain her original statement that killing the five patients to save the one would be morally right, the doctor could explain that since killing the five patients to save the one would likely maximise happiness, that action has one of the most central and paradigmatic hallmarks of rightness. However, since (on my version of Naturalism) that is neither necessary nor sufficient for the action to have any normative property (such as rightness or goodness), the doctor could reason that the action may have other morally relevant natural properties, which make the normative status of the action murkier – such as the property that performing the action would mean a patient would be sacrificed without his consent.

By reasoning this way, the doctor could preserve the Naturalist insight that there is a conceptual connection between an action maximising happiness and its being good, or right, without going too far and simply equating these at a conceptual level, or claiming that saying “action X is good” expresses the same fact as “action X maximises happiness”. As such, the doctor’s normative belief would emerge intact, for she will now acknowledge that an act could fail to be right even if it does maximise happiness.

There would be another upside: reasoning this way about the meaning of the normative language pertinent to this case would, as we have seen, lead the doctor to think about the case in a more nuanced way, helping her understand the moral complexity of the case at hand. This helps quell one worry that comes up when reflecting on Parfit’s doctor case, which is that from Parfit’s commentary it can seem as though the Naturalist is committed to a bizarre way of thinking about morality. But my form of Naturalism, I have shown, has no such bizarre implications.

On my view, to fully understand evaluative claims, we must know which descriptive claims are conceptually tied up with the evaluative statement in question. Our primary example of how this conceptual tie works comes from Meynell (1971): if one does not know that morally good things *generally* produce happiness, then this calls into doubt whether one fully understands what it is for something to be “morally good”, *even though* it is possible for something to be morally good but not produce happiness. This example shows how evaluative and descriptive claims conceptually tie in together, part and parcel. So, this novel form of Naturalism does not aim to distil language into thin evaluative and descriptive forms, but rather elucidates evaluative language by showing how it is conceptually interwoven with descriptive language.

Part iv: Differentiating My Naturalism from Jackson’s Naturalism

Jackson (1998) offers a Naturalist view which seems similar to my own Naturalism, at least at first glance. To clarify my view and convince the reader that my view is distinct from Jackson’s, I will now set out the differences between the two forms of Naturalism and why I think my view is more justified.

As with my view, Jackson avoids claiming that a moral property such as “goodness” is identical with any one natural property, like some simple forms of Naturalism do. Instead, Jackson analyses the meaning of moral properties like “goodness” in terms of the natural properties which feature in platitudes concerning goodness. As Jackson puts it, “we need to identify rightness as the property that satisfies, or near enough satisfies, the folk theory of rightness – and likewise for other moral properties” (1998: 118). By “folk theory”, Jackson means what ordinary speakers would agree on about what is right, and this is why he gives platitudes concerning goodness and rightness a key role in his account: these represent not what particular individuals think about what is right, but what the community of speakers as whole agrees on about what is right (1998: 118). The folk theory of morality, Jackson says, is “the network of moral opinions, intuitions, principles and concepts whose mastery is part and parcel of having a sense of what is right and wrong, and of being able to engage

in a meaningful debate about what is to be done” (1998: 130). This is a point of overlap between Jackson’s view and my own.

On this basis, however, Jackson proposes that a sentence like “X is right” can be analysed down into a “purely descriptive” (1998: 123) sentence of the form “X is P, or X is Q, or X is R...”, where P, Q, R, etc. are “descriptive”, natural properties which “folk morality” associates with rightness. So according to Jackson, “ethical properties are descriptive properties” (1998: 123). It is only that, on Jackson’s view, calling an action *right* does not necessarily mean that it has any one particular natural property that we associate with rightness, such as “maximising happiness”, but rather means that the action has some natural property from a list of (potentially very long) natural properties which are commonly associated with right acts in folk morality.

Jackson’s view, then, is a *reductionist* view according to the definition of “reduction” I appealed to, and explained, above (see Klocksiem 2019). Following Klocksiem, I hold that a meta-ethical view is reductionist if it says that moral claims, which appear to ascribe *normative* properties, really ascribe *descriptive* or *naturalistic* properties. While Jackson does not argue that standard moral claims like “X is good” ascribe any *single* non-normative natural property, he still argues that “X is good” ascribes some *combination* of non-normative natural properties, and that the meaning of “X is good” can be wholly captured by substituting in a comprehensive appropriate list of the relevant non-normative properties.

While Jackson (correctly, in my view) includes a broader range of natural properties in his analysis of normative words than most Naturalists – namely, all those which are part of the “folk theory of morality” – he is therefore a reductionist, nevertheless. As I have explained, reductionist versions of Naturalism cease to make sense if we deny the Dichotomous Claim, since they appeal to the notion of semantically thin descriptive language (which is what they aim to “reduce” moral language to).

Moreover, as I have argued, reductionists are wrong about the relationship between natural properties such as “maximises happiness” and normative properties like “goodness”. Goodness is not substitutable, on my view, by any set of natural properties (e.g. “maximises happiness”, “is what one would desire, if fully informed”, etc.), just as the property of “being a table” is not substitutable by any list of natural properties e.g. “has a flat table”, “has legs”, etc.). However, the property of goodness can fruitfully *be explained in terms of* such natural properties. To better understand the difference between my view and Jackson’s, consider the different ways in which they respond to Parfit’s “Normativity Problem”. To go back to Parfit’s case of the doctor before the Ethics Committee, on Jackson’s view the doctor’s moral belief really *does* amount to a belief about non-normative natural facts – albeit, perhaps, a complicated disjunctive belief about natural facts. This means that Jackson’s view is still vulnerable to Parfit’s objection that, if the doctor follows Jackson’s view, it is as though she has *no normative belief at all*, since (in her own eyes) her moral view really amounts, not to something normative, but to a non-normative belief about natural facts.

By contrast, my view does not face this problem. Unlike Jackson and other reductionists, I do not consider beliefs about natural facts to be (entirely) “non-normative”, since at a semantic level I do not accept the distinction between “thin”, non-normative descriptive language and “thick” or “thin” normative language. Jackson openly admits that he argues that “ethical properties are descriptive properties in the sense of properties ascribed by language that falls on the descriptive side of the famous is-ought divide” (1998: 120). As I have argued above, though, we should not accept the “is-ought divide”, in the sense I have defined that divide.¹¹⁴

Sensing controversy about his reliance on the “is-ought divide”, Jackson proposes a “play-safe” strategy (1998: 120) on which, if any doubt is raised about whether any given word should be classified as “purely descriptive” (or “thinly descriptive”, in other words) we “take it off the list of purely descriptive” (1998: 120). But while Jackson may be prepared to move any given single word off the “purely descriptive” list, he still seems to tacitly

¹¹⁴ See Chapter 2, Section 1 for my definition of the Fact/Value Dichotomy and Chapter 3, Section 2 for my argument against the Fact/Value Dichotomy.

accept the possibility of “pure” descriptive language, which is further evidence that he, unlike the present author, accepts the Dichotomous Claim. Even more clearly, while Jackson says he does not accept that there is a “*sharp* semantic divide between ethical and descriptive terms” he admits he has to “regard the purely descriptive terms as essentially given by a big list of terms that would generally be classified as such”, much like Parfit (2011) (see Chapter 2, Section 1).

Jackson also writes that “even if you belong to the party that thinks that the division between ethical and descriptive vocabulary is a hopeless confusion”, we can still sensibly ask whether what are generally considered moral properties supervene on what are generally considered non-normative natural properties (1998: 121). But this question about supervenience is not pertinent to my focus in this thesis, and this remark from Jackson does nothing to lessen the dependence of his position on the Dichotomous Claim, a Claim which I have already argued against. Moreover, as I argued a moment ago, Jackson’s form of meta-ethical Naturalism continues to suffer problems at the hands of Parfit’s Normativity Problem, whereas mine does not. I have hence now established both why my own view differs from Jackson’s, and why mine is to be preferred to his.

Part v: Anscombe/MacIntyre Counterexamples Again

At the end of Chapter 2, Section 1, I discussed Anscombe’s (1958) and MacIntyre’s (1981) purported counterexamples to the Fact/Value Dichotomy, such as the following:

P1) This watch is grossly inaccurate and irregular in time-keeping.

C) This watch is a bad watch (see MacIntyre 1981: 58 and the discussion in Chapter 1, Section 4 and Chapter 2, Section 1, above).

This example argument seemed to pose a problem for the Fact/Value Dichotomy because it conflicts with Johnston's (1989) thesis that descriptive and evaluative language are strictly logically separate, and that therefore descriptive premises alone cannot entail evaluative conclusions. As I discussed in Chapter 1, Section 4 and Chapter 2, Section 1, however, advocates of the Fact/Value Dichotomy would want to deal with such cases by drawing out a "suppressed premise":

P1) This watch is grossly inaccurate and irregular in time-keeping.

P2) If a watch is grossly inaccurate and irregular in time-keeping, it is a bad watch.

C) This watch is a bad watch.

Since, Johnston would say, we now see that the premises include an evaluative claim (P2), the example is now no longer a threat to the idea that descriptive premises alone cannot entail evaluative conclusions. In reply, I suggested that the viability of this response depends on whether P2 is conceptually independent from P1. If P2 unpacks the meaning of P1 in much the same way as "all bachelors are unmarried" unpacks the meaning of "James is a bachelor", then P2 is not conceptually separate from P1 and thus the two claims could not be viewed as logically separate and different in kind, as Johnston wants to do (so I argued in Chapter 2, Section 1).

The advocate of the Dichotomous Claim would, I said, reply by accepting that, in at least some cases, the concepts mentioned in these Anscombe/MacIntyre cases will be semantically "thick". So, for instance, perhaps "watch" is after all a thick concept, on the grounds that it is part of the concept "watch" that a good watch keeps time accurately. If this is so, then P2 simply unpacks part of the meaning of P1 and the two premises are not logically separate. But in this case, Anscombe/MacIntyre arguments would not threaten the Dichotomous Claim, since the advocate of the Dichotomous Claim can view the case of thick concepts (such as, perhaps, "watch", and arguments featuring that concept) as a

special case. If, the advocate of the Dichotomous Claim will say, we had only semantically thin descriptive premises, these could not entail an evaluative conclusion, since in that case there would be no evaluative “material” in the premises with which to “magic up” an evaluative conclusion.

In light of the arguments of this thesis, we should not accept this response to Anscombe/MacIntyre cases. Since I do not think advocates of the Dichotomous Claim have adequately explained their notion of “thin” descriptive language, we should not view the case of “thick” concepts, and arguments featuring them, as “special” cases at all. Rather, we should, at least provisionally – until or unless the Dichotomous Claim makes good on its notion of thin descriptive language - accept Kovesi’s (1967) view that concepts *quite generally* behave in the way that advocates of the Dichotomous Claim think is exclusive to thick concepts.

As MacIntyre observed, it makes sense to say that a descriptive premise entails an evaluative conclusion in cases where one or more of the concepts in the descriptive premise is “functional” (see Chapter 1, Section 4 and Chapter 2, Section 1, above). In the case of the watch-argument, for example, the fact that it is implicit in the concept “watch” that a good watch tells time accurately allows us to draw an evaluative conclusion out of the seemingly purely descriptive premise. What we have learned from Kovesi (see Chapter 3, Section 2, above) is that understanding the meaning of a word seems always to involve understanding its purpose, even with the case of apparently mundane or “thin” words like “square”, “red”, etc. So in our watch-argument, above, P2 represents the knowledge of the purpose of “watch”, which is part and parcel of a complete understanding of the word “watch”. Kovesi’s insight is that this is not a special case for “thick” concepts only, but rather a necessary part of understanding the meaning of a concept in general – or at least, we can only suppose this to be the case, unless or until the notion of semantically “thin” descriptive language is adequately explained. So again, the “functional” case MacIntyre has in mind should not be regarded as a special case, but rather as the norm, and the default model, until or unless it is shown otherwise.

To link back to my discussion of MacIntyre's views in Chapter 1, Section 1, this is why I consider MacIntyre to be right that evaluative moral claims "just are a kind of factual statement" (1981: 173). But as I remarked in that Section, we needed to work out the specific sense in which this is true. Having rejected the Dichotomous Claim and responded to Parfit's Normativity Problem, I am now in a position to make this clarification. Moral claims are not "just a kind of factual statement" in the sense that they are *reducible* to factual statements, contrary to the claims of reductive Naturalists (see parts iii) and iv) of this Section). Nor can the truths which true moral claims express be captured by any "thin" descriptive factual statement, as Non-Analytical Naturalists believe, since (as I have argued) the notion of semantically "thin" descriptive language is problematic.

Rather, I argue, evaluative moral claims share a kind with factual descriptive claims because *the notion that there are two semantically different kinds of language here is mistaken*. Since we cannot yet make sense of the Johnstonian thesis that descriptive claims are of one semantic kind while evaluative (including moral) claims are of a distinct, separate semantic kind, we should hence pivot to the opposite, MacIntyrean "extreme" of collapsing the Fact/Value Dichotomy entirely, in line with the Classical Teleological understanding of moral language and arguments I discussed in Chapter 1, Section 4.

Moreover, since we ought to reject the Fact/Value Dichotomy, we ought to reject not only Johnston's radical theses, but also other, more moderate meta-ethical views which trade in the Fact/Value Dichotomy, like those forms of Naturalist and Non-Naturalist realism I discussed near the beginning of this Section. As we have seen, many of these views take the idea of semantically thin descriptive language for granted. I have argued that taking this confused idea for granted hinders our understanding of moral language and its relationship with descriptive language. Those philosophers who avoid doing so, like Kovesi (1967) and Meynell (1971), offer more fruitful ways forward, which I have started to develop in this Section. My primary reasons for preferring meta-ethical Naturalism without the Dichotomous Claim, other than the fact that I think the Dichotomous Claim is conceptually confused, are that i) adopting this Non-Dichotomous Naturalism allows us to solve Parfit's Normativity Problem and ii) the Non-Dichotomous Naturalism fits with, and

explains, the linguistic data (Meynell's Observation and the Converse Meynell Observation, see above).

Thesis Conclusion – Meta-Ethical Naturalism and the Fact/Value Dichotomy

The primary aim of this thesis has been to clarify the nature of the Fact/Value Dichotomy, and then to undermine that Dichotomy. To achieve this, I scrutinised Johnston's (1989) defence of the Fact/Value Dichotomy in Chapter 1. I found problems with Johnston's defence of the Dichotomy: first, it appeals to the untenable idea that we can conduct meta-ethical inquiry in an entirely ethically neutral way (Chapter 1, Section 3) and, second, its viability turns out to depend on Classical Teleological normative ethical views being false, which I argued is problematic. By identifying these problems with Johnston's position, which have not been correctly identified and explained in previous literature, I reduced the viability of Johnston's way of defending "Emotivism" and the Fact/Value Dichotomy.

While these problems make Johnston's position significantly less attractive, I explained that the Fact/Value Dichotomy can survive despite these problems. I found that not all advocates of the Fact/Value Dichotomy need to buy in to the radical form of the Dichotomy which Johnston defends. I outlined the "Humean Claim" endorsed by Johnston, which says that all claims which are either descriptive or evaluative must be one or the other and cannot be both. It also says that sets of only descriptive premises cannot entail evaluative conclusions. I then contrasted this stronger version of the Fact/Value Dichotomy with the Dichotomous Claim, which is endorsed by, for example, Parfit (2011). This weaker version of the Dichotomy allows that there can be "thick" claims, which are both descriptive and evaluative, but maintains that "thin" (or "pure") descriptive and "thin" evaluative claims are possible too. I also clarified that both the Humean Claim and the Dichotomous Claim are *semantic* theses: they allege that "thin" descriptive claims have a fundamentally different kind of meaning to that of "thin" evaluative claims, such that there is a genuine semantic Dichotomy between two kinds of meaning here.

At first sight, the Dichotomous Claim seems both more plausible than the Humean Claim, and, indeed, highly plausible in its own right. It appears to be well-supported by ordinary

language, and it gives a more natural explanation of the apparent phenomenon of “thick” concepts than, for instance, the Humean Claim does.

To give fair scrutiny of the Dichotomous Claim, I examined both Putnam’s (2002) and Lovibond’s (1983) criticisms of the Fact/Value Dichotomy. I found Putnam’s arguments to weigh against only the Humean Claim, and not the Dichotomous Claim, while Lovibond’s arguments, while they speak against both forms of the Fact/Value Dichotomy, depend on an onerous kind of essentialism. This led me to the conclusion that, to make a more potent argument against the Fact/Value Dichotomy, we would need to avoid the kind of essentialism which Lovibond appeals to.

In Chapter 3, Section 1, I identified a way of doing this, drawing from Mulhall’s (2007) interpretation of Wittgenstein on private language. The basic idea is, rather than to argue that all language is essentially “thick” (like Lovibond), we should instead probe whether the idea of semantically “thin” language makes sense, much like how Wittgenstein (according to Mulhall) probes whether the idea of logically “private” language makes sense. This approach, I argue, does not rely on any kind of essentialism, giving it an advantage over Lovibond’s strategy.

Moreover, as I show in Chapter 3, Section 2, this strategy can be applied to great effect against the Fact/Value Dichotomy. Drawing from Kovesi (1967), I argued that the impression that some words and claims are “thin” is misleading. The notion that we can have meaningful “thin” descriptive language turns out to depend on the idea that we can understand the meaning of a word independently of understanding the purpose of the word, but I made the case that that idea is conceptually confused.

What I claim follows from my argument is not so much that the Dichotomous Claim is false, but rather that its notion of semantically thin descriptive language has not yet been satisfactorily explained. The impression that some of our extant language is “thin” descriptive language is illusory, I have argued. This is consistent with the possibility that

the advocate of the Dichotomous Claim, in future, may somehow provide a new and satisfactory explanation of their idea. Unless or until such an explanation is provided, however, embracing the Dichotomous Claim is untenable. My argument against the Dichotomous Claim also meets the desiderata laid out at the end of Chapter 2, and hence, in my view, marks, at the very least, an improvement upon previous arguments which aimed at the same target (e.g. Lovibond 1983).

In Chapter 3, Section 3, I began exploring the implications of my rejection of the Dichotomous Claim for contemporary meta-ethics, focusing especially on meta-ethical Naturalism. I argued that Parfit's "Normativity Problem" only applies to versions of meta-ethical Naturalism which accept the Dichotomous Claim. While the Normativity Problem does, I think, give good reason to reject the forms of Naturalism to which it applies, I outline a different form of Naturalism which handles the problem convincingly, and in any case gives a good explanation of the relevant linguistic data. I therefore conclude that the non-Dichotomous version of meta-ethical Naturalism which I advanced in Chapter 3, Section 3, is plausible, and preferable to versions of realism (both Naturalist and Non-Naturalist) which embrace the Dichotomous Claim.

I now want to relate my findings back to the Johnston-MacIntyre debate that I discussed in Chapter 1, Section 1. To recap, I considered the "Emotivist" view, on which the rational adjudication of moral disagreements is impossible. Taking Johnston's version of "Emotivism" as an example in Chapter 1, Section 2, we found that Johnston denies the possibility of the rational adjudication of moral disagreements due to his commitment to the Fact/Value Dichotomy. In keeping with this Dichotomy, Johnston understands disagreements about matters of non-normative fact to be entirely different in kind from disagreements about moral questions. In the former, he claims, there is the possibility of rationally resolving disagreement through appeal to evidence, empirical investigation, and so forth. In the latter case, however, no rational adjudication is achievable, because (Johnston says), when it comes to moral claims, we are dealing with mere assertions and counter-assertions about what ought to be done, and since the persons making these claims are not making any claims about reality that we might be able to investigate or find

evidence for, there can be no question of rationally supporting or criticising any given moral claim.

This reminds us of how, in Chapter 1, Section 1, we discussed how Johnston's defence of the Fact/Value Dichotomy supports "Emotivism", by apparently giving a principled explanation of why the rational adjudication of moral disagreements is impossible (rather than simply challenging or difficult, as optimistic realists like Enoch (2011) maintain).

To undermine this Johnstonian way of defending "Emotivism", I have targeted the Fact/Value Dichotomy on which Johnston's position depends. It turns out, I have argued, that Johnston's view that there is a "logical gulf" (Meynell 1971) between descriptive and evaluative claims is mistaken: we have no reason to believe that, when making a moral claim, we are saying nothing about reality, nothing that can be investigated or supported by evidence. Rather, I argued, to call an action "good" (for example) is, at least in a standard case to indicate that one takes that action to possess some properties from a range of relevant natural properties (e.g. "produces happiness", "does the revealed will of God", etc.), much in the same way as calling something a "table" indicates that one takes the object one is talking about to possess some properties from a range of relevant natural properties ("has legs", "has a flat surface", etc.) But crucially, I have argued, contrary to reductive Naturalism, saying that an action is "good" is not *reducible* to saying it has some such natural properties. We can *partially explain* our meaning in calling something "good" by stating the natural properties which we are thereby attributing to it, but doing so will never give an exhaustive explanation of our meaning. For, as Kovesi (1967) teaches us, giving a list of the natural properties which are logically related to terms like "good" and "table" never gives a full picture of the meaning of those words, since to fully understand the word, we also need to understand its purpose.

So while my investigation into the Fact/Value Dichotomy was initially targeted at Johnston's view, we have ended up not only disagreeing with Johnston, but also with all meta-ethical positions which rest upon the Fact/Value Dichotomy in either its stronger or weaker form.

In terms of the implications of my non-Dichotomous version of meta-ethical Naturalism for moral disagreement, I have not had space to address this issue in detail. What exactly would follow from my view for how it might be possible to rationally resolve moral disagreements is a matter for future research. However, I have, at least, undermined one major source of pessimism about the possibility of rationally resolving moral disagreements: the Fact/Value Dichotomy. I believe we hence have good reason to, like Enoch (2011), hold out hope for the possibility of rationally resolving moral disagreements.

Bibliography

Agam-Segal, Reshef, and Edmund Dain (eds.) (2017) *Wittgenstein's Moral Thought*, London: Routledge.

Anscombe, G.E.M. (1958) "Modern Moral Philosophy", *Philosophy* vol. 33 (124), pp. 1-19.

Aquinas, Thomas (1485) *Summa Theologica*, www.newadvent.org/summa/index.html, accessed 13th March 2019.

Aristotle (2000) [340 BC], *Nicomachean Ethics*, ed. and trans. Roger Crisp, Cambridge: Cambridge University Press.

Ayer, Alfred Jules (1952) *Language, Truth and Logic*, London: Victor Gollancz.

Ayer, Alfred Jules (1972) *Philosophical Essays*, Basingstoke: Palgrave-Macmillan.

Blackburn, Simon (1984) *Spreading the Word*, Oxford: Clarendon.

Blackburn, Simon (1998) *Ruling Passions*, Oxford: Clarendon.

Brink, David O. (1986) "Externalist Moral Realism", *Southern Journal of Philosophy* vol. 24 (Supplement): 23–40.

Brink, David O. (1989) *Moral Realism and the Foundations of Ethics*, Cambridge: Cambridge University Press.

Burton, Stephan L., (1992) "'Thick' Concepts Revised", *Analysis* vol. 52(1), pp. 28–32.

Canfield, John V. (2001) "Private Language: The Diary Case", *Australasian Journal of Philosophy* vol. 79 (3), pp. 377-394.

Cavell, Stanley (1979) *The Claim of Reason*, New York: Oxford University Press.

Clark, Kelly James (ed.) (2016) *The Blackwell Companion to Naturalism*, Oxford: Blackwell.

Chappell, Timothy (2013) "There Are No Thick Concepts", in Kirchin (ed.)

Christensen, Anne-Marie Søndergaard (2017), " "What is Ethical Cannot Be Taught" – Moral Theories as Descriptions of Grammar", in Agam-Segal and Dain (eds.)

Copp, David and David Zimmerman (eds.) (1985) *Morality, Reason, and Truth*, Totowa, NJ: Rowman and Allanheld.

Craig, Edward (1990) *Knowledge and the State of Nature*, Oxford: Oxford University Press.

Crary, Alice and Rupert Read (eds.) (2000) *The New Wittgenstein*, London: Routledge.

Dancy, Jonathan (1995) "In Defence of Thick Concepts", *Midwest Studies in Philosophy* vol. 20, pp. 263-279.

Dancy, Jonathan (2013) "Practical Concepts", in Kirchin (ed.)

De Mesel, Benjamin (2015) "Wittgenstein, Meta-Ethics and the Subject Matter of Moral Philosophy", *Philosophical Perspectives* 22, pp. 69-98.

Diamond, Cora (1996) "Wittgenstein, Mathematics and Ethics", in Sluga and Stern (eds.)

Diamond, Cora (2000) "Ethics, Imagination, and the Method of Wittgenstein's *Tractatus*", in Crary and Read (eds.)

Dreier, Jamie (2004) “Meta-Ethics and the Problem of Creeping Minimalism”, *Philosophical Perspectives* vol. 18 (1), pp. 23-44.

Dummett, Michael (1973) *Frege: Philosophy of Language*, London: Duckworth.

Dummett, Michael (2006) *Thought and Reality*, Oxford: Oxford University Press.

Eklund, Matti (2017) *Choosing Normative Concepts*, Oxford: Oxford University Press.

Elstein, Daniel Y. and Thomas Hurka (2009) “From Thick to Thin: Two Moral Reduction Plans”, *Canadian Journal of Philosophy* vol. 39(4), pp. 515–36.

Enoch, David (2011) “How is Moral Disagreement a Problem for Moral Realism?”, *The Journal of Ethics* vol. 13 (1), pp. 15-50.

Fehige, Christoph and Ulla Wessels (eds.) (1998) *Preferences*, Berlin: de Gruyter.

Fogelin, Robert (1996) *Wittgenstein*, London: Routledge.

Fricke, Miranda (2016) “What’s the Point of Blame? A Paradigm-Based Explanation”, *Nous* vol. 50 (1), pp. 165-183.

Geertz, Clifford (1973) *The Interpretation of Cultures*, New York: Basic Books.

Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Pragmatic Justification of Morality*, Cambridge, MA: Harvard University Press.

Gibbard, Allan (1998) “Preference and Preferability”, in Fehige and Wessels (eds.)

Gibbard, Allan (2003) *Thinking How to Live*, Cambridge, MA: Harvard University Press.

Gibbard, Allan (2009) "A Pragmatic Justification of Morality", in Voorhoeve (2009).

Ginsborg, Hannah (2018) "Normativity and Concepts", in Star (ed.)

Gowans, Christopher (2000a) "Introduction", in Gowans (ed.)

Gowans, Christopher (ed.) (2000b) *Moral Disagreements: Classic and Contemporary Readings*, London and New York: Routledge.

Greco, John, and David Henderson (eds.) (2015) *Epistemic Evaluation: Purposeful Epistemology*, Oxford: Oxford University Press.

Gregory, Alex (2017) "Are All Normative Judgments Desire-Like?" *Journal of Ethics and Social Philosophy* vol. 12 (1), pp. 29-55.

Harcourt, Edward, and Alan Thomas (2013) "Thick Concepts, Analysis and Reductionism", in Kirchin (ed.)

Harcourt, Edward (2019) 'Moral Concepts, "Natural Facts" and Naturalism: Outline of a Wittgensteinian Moral Philosophy', in Kuusela and De Mesel (eds.)

Hare, John E. (2001) *God's Call: Moral Realism, God's Commands, and Human Autonomy*, London: Eerdmans.

Harrison, Bernard (2012) "Kovesi's Refutation of Hume", in Tapper and Mooney (eds.)

Heal, Jane (2013) "Social Anti-Individualism, Co-Cognitivism, and Second-Person Authority", *Mind* vol. 122 (486), pp. 339-371.

Hicks, R.D. (1910) *Stoic and Epicurean*, London: Longmans Green.

Hieronymi, Pamela (2005) 'The Wrong Kind of Reason', *The Journal of Philosophy* 102, pp. 437-457.

Hume, David (1985) [1739] *A Treatise of Human Nature*, London: Penguin.

Jackson, Frank (1998) *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Oxford University Press.

Johnston, Paul (1989) *Wittgenstein and Moral Philosophy*, London: Routledge.

Kant, Immanuel (2012) [1785] *Groundwork for the Metaphysics of Morals*, ed. Mary Gregor and Jens Timmermans, Cambridge: Cambridge University Press.

Kovesi, Julius (1967) *Moral Notions*, London: Routledge & Kegan Paul.

Kenny, Anthony (2001) *Essays on the Aristotelian Tradition*, Oxford: Clarendon.

Kirchin, Simon (ed.) (2013) *Thick Concepts*, Oxford: Oxford University Press.

Klocksiam, Justin (2019), "Against Reductive Ethical Naturalism", *Philosophical Studies* vol. 176, pp. 1991-2010.

Kripke, Naming and Necessity (1980) *Naming and Necessity*, Cambridge, MA: Harvard University Press.

Kuusela, Oskari, and Benjamin De Mesel (eds.) (2019) *Ethics in the Wake of Wittgenstein*, London: Routledge.

Lovibond, Sabina (1983) *Realism and Imagination in Ethics*, Oxford: Basil Blackwell.

Lovibond, Sabina (2002) *Ethical Formation*, Cambridge, MA: Harvard University Press.

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

MacIntyre, Alasdair (1959) "Hume on "Is" and "Ought"", *The Philosophical Review* vol. 68 (4), pp. 451-468.

MacIntyre, Alasdair (1967) *A Short History of Ethics*, London: Routledge.

MacIntyre, Alasdair (1981) *After Virtue*, London: Duckworth.

MacIntyre, Alasdair (1988) *Whose Justice? Which Rationality?*, London: Duckworth.

MacIntyre, Alasdair (2016) *Ethics in the Conflict of Modernity*, Cambridge: Cambridge University Press.

Marcus, Ruth Barcan (1967) "Essentialism in Modal Logic", *Noûs* vol. 1: pp. 91–96.

McDowell, John (2000) "Non-Cognitivism and Rule-Following", in Crary and Read (eds.)

McDowell, John (2009) *The Engaged Intellect: Philosophical Essays*, Cambridge, MA: Harvard University Press.

McManus, Denis (2013) "Heidegger, Wittgenstein and St. Paul on the Last Judgment: On the Roots and Significance of the "Theoretical Attitude"", *British Journal for the History of Philosophy* vol. 21 (1), pp. 143-164.

McPherson, Tristram, and David Plunkett (eds.) (2017) *The Routledge Handbook of Metaethics*, London: Routledge.

Meynell, Hugo (1971) "The Objectivity of Value Judgments", *The Philosophical Quarterly* vol. 21 (83), pp. 118-131.

Moore, G.E. (1903) *Principia Ethica*, Cambridge: Cambridge University Press.

Mulhall, Stephen (2002) "Ethics in the Light of Wittgenstein", *Philosophical Perspectives* vol. 31 (3), pp. 293-321.

Mulhall, Stephen (2007) *Wittgenstein's Private Language*, Oxford: Oxford University Press.

Parfit, Derek (1984) *Reasons and Persons*, Oxford: Oxford University Press.

Parfit, Derek (2011) *On What Matters*, vol. 2, Oxford: Oxford University Press.

Plato (2007) [381 BC] *Republic*, trans. Melissa Lane, London: Penguin.

Putnam, Hilary (1981) *Reason, Truth and History*, Cambridge: Cambridge University Press.

Putnam, Hilary (2002) *The Collapse of the Fact/Value Dichotomy and Other Essays*, Cambridge, MA: Harvard University Press.

Queloz, Matthieu (2019), "The Points of Concepts: Their Types, Tensions and Connections", *Canadian Journal of Philosophy* vol. 49 (8), pp. 1122-1145.

Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen (2004) "The Strike of the Demon: Fitting Pro-Attitudes and Value", *Ethics* vol. 114 (3), pp. 391-423.

Russell, Bertrand (1905) "On Denoting", *Mind* vol. 14 (56), pp. 479-493.

Russell, Bertrand (1918) *Mysticism and Logic*, London: Longmans Green.

Sayre-McCord, Geoffrey (2012) "Moral Realism", in *The Stanford Encyclopedia of Philosophy*, edited by Edward Zalta (October 2005), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/moral-realism/> [accessed 28th July 2017].

Sayre-McCord, Geoffrey (2012) "Metaethics", in *Stanford Encyclopedia of Philosophy* (ed. Edward Zalta), <http://plato.stanford.edu/archives/spr2012/entries/metaethics/> [accessed 21 January 2015].

Schoenbaumsfeld, Genia (2016) *The Illusion of Doubt*, Oxford: Oxford University Press.

Schroeder, Mark (2008) *Being For: Evaluating the Semantic Programme of Expressivism*, Oxford: Oxford University Press.

Schroeder, Mark (2017) 'Normative Ethics and Meta-Ethics', in McPherson and David Plunkett (eds.)

Sluga, Hans, and David Stern (eds.) (1996) *The Cambridge Companion to Wittgenstein*, Cambridge: Cambridge University Press.

Smith, Michael A. (1994) *The Moral Problem*, Oxford: Wiley-Blackwell.

Star, Daniel (ed.) (2018) *The Oxford Handbook of Reasons and Normativity*, Oxford: Oxford University Press.

Stevenson, Charles L. (1945) *Ethics and Language*, New Haven: Yale University Press.

Strawson, P.F. (1959) *Individuals*, London: Methuen.

Streumer, Bart (2017) *Unbelievable Errors: An Error Theory About All Normative Judgments*, Oxford: Oxford University Press.

Sturgeon, Nicholas (1985) "Moral Explanations", in Copp and Zimmerman (eds.)

Suikannen, Jussi (2016) "Naturalism in Metaethics", in Clark (ed.), pp. 351-368.

Tapper, Alan, and T. Brian Mooney (eds.) (2012) *Meaning and Morality: Essays on the Philosophy of Julius Kovesi*, Leiden: Brill Publishers.

Tappolet, Christine (2004) "Through Thick and Thin: Good and its Determinates", *Dialectica* vol. 58(2), pp. 207–21.

Williams, Bernard (1985) *Ethics and the Limits of Philosophy*, London: Fontana.

Wittgenstein, Ludwig (1953) *Philosophical Investigations*, trans. G.E.M. Anscombe, Oxford: Blackwell.

Wittgenstein, Ludwig (1993) *Philosophical Occasions, 1912-51*, ed. James C. Klagge & Alfred Nordmann, Indianapolis: Hackett.

Vayrynen, Pekka (2013) *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*, Oxford: Oxford University Press.

Vayrynen, Pekka (2017) "Thick Ethical Concepts", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.).

Voorhoeve, Alex (ed.) (2009) *Conversations on Ethics*, Oxford: Oxford University Press.

