



LUND UNIVERSITY

Value Grounded on Attitudes. Subjectivism in Value Theory

Fritzson, Fritz-Anton

2014

[Link to publication](#)

Citation for published version (APA):

Fritzson, F.-A. (2014). *Value Grounded on Attitudes. Subjectivism in Value Theory*. [Doctoral Thesis (monograph), Department of Philosophy]. Lund University (Media-Tryck).

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Value Grounded on Attitudes

Subjectivism in Value Theory

Fritz-Anton Fritzon



LUND
UNIVERSITY

© Fritz-Anton Fritzon

Faculty of Humanities, Department of Philosophy
ISBN 978-91-7473-993-0 (printed publication)
ISBN 978-91-7473-994-7 (electronic publication)

Printed in Sweden by Media-Tryck, Lund University
Lund 2014

Cover illustration by Elise Catjar



Contents

Preface	7
Acknowledgements	9
1. COMMON GROUND: WHAT SUBJECTIVISTS AND OBJECTIVISTS (MAY) SHARE	11
I. Introduction	11
II. First and Second Order Projects	12
III. Different Second Order Projects	15
IV. Supervenience	19
V. Final and Intrinsic Value	22
VI. Value Bearers	25
VII. Value and Normative Practical Reasons	31
VIII. Good and Good For	37
IX. A Brief Roadmap	41
2. VALUE GROUNDED ON ATTITUDES: SUBJECTIVISM IN VALUE THEORY	45
I. Introduction	45
II. Backgrounding Attitudes	46
III. Two Interpretations of Preferentialism in Value Theory	52
IV. Ground and Object	57
V. How Subjectivists Can Analyse Intrinsic (Final) Value	60
VI. Objections	62
VII. First-Person and Third-Person Perspectives	65
VIII. A Bleak View on which Nothing Matters?	68

3. THE IMPORTANCE OF BEING REAL: VALUE AND HYPOTHETICAL ATTITUDES	71
I. Introduction	71
II. What If There Were No Subjects with Attitudes?	72
III. What If We Had 'Bad' Attitudes?	74
IV. What If Our Attitudes Are Not as We Suppose?	78
V. Invariance and Consistency	80
VI. Is Value-Constitution a Conscious Act?	92
4. WHICH ATTITUDES? MOTIVATION AND IDEALIZATION	97
I. Introduction	97
II. Values Are Not Grounded on Beliefs	100
III. Motivational Internalisms	102
IV. Idealization	113
V. Restrictions on Value-Grounding Attitudes?	119
5. WHOSE ATTITUDES? INDIVIDUALIST RELATIVISM AND UNIVERSALIST ABSOLUTISM	129
I. Introduction	129
II. The Spectator and the Participant	131
III. Relativism and Absolutism	134
IV. Implications of Relativist and Absolutist Subjectivism	142
V. What Kind of Relativism Is This?	147
VI. A Presumption in Favour of Individualist Relativism	153
VII. Is This View Objectionably Self-Regarding?	158
6. OH MY GOODNESS! SUBJECTIVISM AND RELATIONAL VALUE	163
I. Introduction	163
II. Goodness-Type Monism and Goodness-Type Dualism	165
III. A Subjectivist Analysis of Goodness-For	172
IV. Implications of the Analysis	175
7. PROJECTING VALUE: MAKING GOOD ON A METAPHOR	179

I.	Introduction	179
II.	Projection	181
III.	Projectivism and the Ontological Status of Value	193
8.	SUBJECTIVISM AND EVALUATIVE LANGUAGE: THE ERRORS OF OUR WAYS	199
I.	Introduction	199
II.	Subjectivism and Naturalism	200
III.	Subjectivism and Non-Cognitivism	203
IV.	Subjectivism and Error Theory	211
	Afterword	217
	Bibliography	223

Preface

It is commonly thought that what separates subjectivists from objectivists with regard to value is that whilst subjectivists take value to be related to the attitudes of subjects, objectivists deny any such relation. But that value is in *some* way or other related to attitudes, even necessarily so, is not something that objectivists need to deny. Perhaps we might be tempted to say instead that on subjectivist views value depends on attitudes whilst objectivists deny this dependence. But this would be even less accurate since there is an important sense of *depend* in which a subjectivist may deny that value necessarily depends on attitudes whilst some objectivists may affirm that values necessarily depend on attitudes in that sense. What it means for value to ‘depend on attitudes’ is open to more than one interpretation, not all of which are best seen as subjectivist. Thus depicted, the dialectical situation could easily lead us to give up altogether on the quest to find a clear demarcation line between subjectivism and objectivism, and to conclude instead that classifying views as either subjectivist or objectivist is not so interesting after all. But this, I believe, would be a mistake. Subjectivism and objectivism really are distinct positions in value theory: there is indeed something characteristic of subjectivist views that all objectivists deny, and this something concerns the distinctive role of attitudes in subjectivist analyses of value. For the subjectivist, I will argue, attitudes play the role of *grounding* value. My central focus will be on how we might understand the way values are grounded on attitudes and the implications of conceiving of value as being so grounded.

In this, I aim to clear the field of some common misunderstandings and offer a sympathetic formulation of value subjectivism. Whilst I do defend the subjectivist view against some common objections, I do not embark on a full-blown defence of value subjectivism. What I do offer is perhaps better understood as a preamble to a more thoroughgoing defence of the view. After all, a defence of any position would have to be preceded by a suitable articulation of the position to be defended. In order to settle

the debate between subjectivism and objectivism, the many objectivist alternatives would need a much more thorough consideration than I give them in these pages. I have opted for a more positive project—focusing on how subjectivism is best understood—rather than on the negative project of trying to undermine objectivism (a project that has been pursued by others). I do hope, however that when some of the confusions have been cleared up and the choice between subjectivism and objectivism lies more clearly before us, subjectivism will appear to be a more attractive position than many philosophers have been inclined to think. Though I doubt that many objectivists will be converted, an objectivist reader will hopefully get a clearer picture of what I take to be their strongest opponent and of how the subjectivist view differs from their own.

Acknowledgements

Philosophy is a team sport and in no possible world could I have done this alone. First of all I would like to thank my supervisors Wlodek Rabinowicz and Toni Rønnow-Rasmussen whose influence on this work is undeniable. They have always been standing by ready to give valuable critiques of my ideas, virtually none of which have thus far survived a meeting with Wlodek and Toni unaffected. I would also like to thank Johan Brännmark who read and commented on the whole manuscript in the role of departmental opponent at my final seminar in February this year. Johan's perspective helped me to clarify my project, though no doubt it could be made clearer still. I stand in debt also to all of the participants of the weekly Research Seminar in Practical Philosophy at Lund University who read and discussed most of the chapters at various stages of their evolution. On one of these occasions, David Alm made a suggestion that went on to play a key role in my project. I am grateful to David for his suggestion and for his subsequent attempts to expose weaknesses of his own idea. Björn Petersson has always been helpful with both philosophical and practical matters. I further stand in debt to the participants of the Doctoral Student Seminar who have read and discussed my writings in even their embryonic stages of development over the years. I have particularly benefitted from the advice of Frits Gåvertsson, Henrik Andersson, Andrés Garcia and Eric Brandstedt. Frits and Henrik have generously provided written comments on the whole manuscript and Andrés Garcia has frequently discussed with me many of the ideas of the thesis (and much else besides) both inside and outside of the seminar room. Some ideas in the thesis were presented at the Understanding Value conference at The University of Sheffield in July 2013 and at the biannual Reading-Lund conference in May 2010. I wish to thank the Division of Practical Philosophy at the Department of Philosophy for generously sending me and a handful of my Doctoral Student colleagues on an inspirational trip to Oxford in October 2012 as well as for financing a helpful course on how to finish a dissertation thesis on time. My colleagues

in Theoretical Philosophy and Cognitive Science as well as the administrative personnel at the department also deserve mention for contributing to the stimulating work environment in Kungshuset were I have been very fortunate to have had a work place from February 2009 to June 2014. This period has really been the best and most exciting of my life thus far. For support on a personal level, I wish to thank my family and my friends: my parents Elis and Marianne Fritzson, and my sister Elise Catjar who created the cover illustration. Finally, I wish to thank Ieva Vasilionyte for all her love and support both personal and philosophical. Ieva, you are of the highest value to me!

Fritz-Anton Fritzson, Lund, July 2014

1. COMMON GROUND: WHAT SUBJECTIVISTS AND OBJECTIVISTS (MAY) SHARE

[I]t is ambition enough to be employed as an under-labourer in clearing the ground a little, and removing some of the rubbish that lies in the way to knowledge.

John Locke, *Essay Concerning Human Understanding*

I. Introduction

In this work, subjectivism and objectivism will be seen as rival analyses of the nature of value. In order for subjectivists and objectivists to be rivals, their respective analyses must be focused on (at least partly) the same thing—they must have a (at least partly) shared *analysandum*. If this turns out not to be the case, then subjectivists and objectivists are only talking past each other and thus not really disagreeing; hence there would be no subjectivist/objectivist debate in value theory to begin with. That subjectivism and objectivism are opposing views will be treated as a guiding assumption to be given up only if faced with strong considerations to the contrary. Unfortunately, the possibility that subjectivists and objectivists are in fact talking about different things cannot be ruled out easily once and for all, especially given that it is not always made clear in the literature what belongs to a particular proposed analysis of value and what belongs to the

(supposedly shared) analysandum. I have tried to make this as clear as I can in my discussion.

This introductory chapter aims to characterize the analysandum that I believe subjectivists and objectivists alike seek to analyse: the common philosophical project in which I take subjectivists and objectivists to be engaged. I identify a number of points on which subjectivists and objectivists may agree. This will make up the common ground against which the differences between the two theories can be characterized. Along the way, I introduce a number of terms and distinctions that will be used in subsequent chapters. Beginning in Chapter 2 I then focus on subjectivism, aiming to identify what is distinctive about the subjectivist viewpoint in value theory.

II. First and Second Order Projects

In saying that subjectivism and objectivism are rival analyses of the nature of value, I mean to distinguish these views from any first order views. A first order view takes an evaluative or normative stance, something that I aim to avoid in this work. Subjectivism and objectivism will here be treated as second order views about value. As such, subjectivists and objectivists both aim to tell us what value *is*, and what it means for something to be valuable, but do not take a stance on what is valuable. Hence, subjectivism and objectivism do not directly aim to tell us what is good and bad, better and worse, beautiful and ugly, what we have reason to strive for, or how one ought to live. Of course, some subjectivists and objectivists have also argued for various first order views. Does this mean that their respective analyses of the nature of value also have evaluative or normative content after all? The second order value theorist could respond in the words of John Cleese's character in a well-known sketch: 'Not necessarily; I could be arguing in my spare time!'

Understood as second order theories, subjectivism and objectivism are here taken to be in and of themselves neutral with respect to first order

(axiological and normative¹) views. Subjectivists and objectivists—though they disagree with regard to the second order issue of the nature of value—might well be in agreement concerning axiological and normative matters.² A subjectivist and an objectivist might, for example, both accept or reject the familiar axiological thesis that happiness is the greatest good or might both adhere to or oppose the equally familiar normative principle that one ought never to treat persons merely as means, but always as ends in themselves.

I am operating under the assumption that it is, in general, important to try to uphold the distinction between first- and second order value inquiries. That said, however, the distinction between first- and second order issues may not always be perfectly clear-cut, and without undertaking a thorough case-by-case investigation one could not conclude that it never breaks down. Even though the subjectivist view that I develop in the following chapters does not in and of itself imply any particular first order view, it does (as will be seen more clearly later on) have some interesting and important first order implications when fed with the relevant empirical facts—that is, when applied to the real world. For example, under the highly plausible empirical assumption that our attitudes have greatly heterogeneous objects, one such implication seems to be axiological pluralism: the view that several different things, such as, for example, knowledge, peace, biodiversity or artistic expression, are valuable for their own sake.³ This kind of pluralism should

¹ Axiological theories aim to tell us what is valuable: what is good and bad, better and worse, beautiful and ugly, etc. A normative theory aims to tell us what is right and wrong, what we have reason to do, how one ought to live, etc. I understand the normative and the axiological to be intimately connected. As will be argued below, that something is valuable implies that there are normative reasons to respond to it in some way or other. Sometimes, second order value theory is called formal axiology, whilst the corresponding label for first order value theory is substantive axiology.

² That subjectivists and objectivists need not disagree about first order matters is stressed by, among others, David Gauthier and Toni Rønnow-Rasmussen. Rønnow-Rasmussen (2003) maintains that ‘subjectivists and objectivists may well share each other’s evaluations, and they need not at all disagree about what objects are the appropriate bearers of value’ (p. 247). Gauthier (1986) similarly insists that subjectivism and objectivism are ‘equally neutral with respect to the nature of the objects of intrinsic value’ (pp. 47–48). Gauthier here uses the term *intrinsic*, but to talk about *final* value is, as I will indicate below, more appropriate. We need not suppose that all final values are intrinsic.

³ This implication is noted in Rabinowicz & Österberg (1996): ‘Our intrinsic preferences seem to have greatly heterogeneous objects. Were we to retain those preferences under

be contrasted with monistic axiological views, such as that of the axiological hedonist who holds that one and only one thing—pleasure—is valuable for its own sake. Another empirical assumption—which can scarcely be questioned—is that attitudes vary from one person to another.⁴ Given the subjectivist idea that values are (in a yet to be specified way) determined by attitudes, this seems to naturally give rise to the view that what is valuable relative to one person need not be valuable relative to another.⁵ Even though this relativistic implication does not by itself force on us any specific normative commitment⁶, it may very well make a difference to the outcomes when certain normative theories are being applied.

The label *subjectivism* has sometimes been applied to first order views. Derek Parfit, for example, calls the view according to which ‘we have most reason to do whatever would best fulfil or achieve our present desires or aims’ a subjectivist view.⁷ This usage is obviously fundamentally different from mine. That something is good or bad, better or worse, or that we have more reason to do one particular thing versus another does not, as I have insisted, belong to an analysis of value but rather to the analysandum—it is

reflection, it would follow that there are several irreducibly distinct intrinsic values’ (pp. 10–11).

⁴ Of course, it remains an interesting and often controversial empirical question just how large the variation in attitudes really is. Few doubt that there is a class of attitudes that we all share in virtue of being human. Several entries on Donald E. Brown’s *List of Human Universals* may, for example, be interpreted as describing universal human attitudes.

⁵ It should be noted that the concept of ‘good relative to’ will be treated as distinct from that of ‘good for’. The idea that something is good or bad *for* someone is familiar to all of us and the expression ‘good for’ is (in some usages) used in everyday language to make first order, evaluative claims (such as ‘studying philosophy is good for you’). ‘Good relative to’, by contrast, is something of a term of art that I will be using to characterize a certain second order thesis.

⁶ Krister Bykvist (1996) argues, to the contrary, that—precisely in virtue of its relativistic implications—a subjectivist theory of value carries a negative normative commitment in that it cannot be combined with a consequentialist normative theory (p. 14). His argument will be discussed in Chapter 5. I maintain there that value subjectivism, despite its relativistic implications, *is* compatible with consequentialist normative theory, as well as with other normative theories.

⁷ Parfit (2011) (p. 1)

part of that which is to be analysed. I am not claiming that usages of the term *subjectivism* (or *objectivism*) that differ from mine are wrong, but I do think that we stand to gain greater clarity if these labels are reserved for second order views. It is therefore with this meaning that I will be using these terms throughout.

III. Different Second Order Projects

Leaving aside first order views, there are other second order projects from which I wish to distinguish the project that I take subjectivists and objectivists to be engaged in. One such project is that of analysing the linguistic meaning of evaluative utterances and determining which type of mental states are expressed by such utterances. This is a different project from that of analysing the nature of value. Since some views of a linguistic or semantic kind have also been labelled *subjectivism*, it is important to distinguish those from subjectivism in my sense.

On one such view, evaluative utterances are used to report the speaker's own attitudes. On a crude version of this view, 'slavery is bad' has the same meaning as 'I [the speaker] disapprove of slavery'. The utterance is a statement about the speaker's attitude that is true if the speaker really does disapprove of slavery and false otherwise. David Hume seems to have given voice to (a slightly more sophisticated version of) this view when he wrote that '[W]hen you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it'.⁸ Thomas Hobbes is another philosopher to whom this type of view is often ascribed.⁹

⁸ *Treatise* (iii.i.i.). As expressed in this passage, Hume's view is that evaluative judgments report the speaker's attitudes, but they do more than just that. They also make claims to the effect that those attitudes are somehow based in the speaker's nature.

⁹ When Stephen Darwall (2000) argues that, contrary to common interpretation, Hobbes was not a subjectivist but instead a projectivist, he is using subjectivism to refer to this view of what evaluative statements mean (pp. 315–317). As will become clear later on, what I will mean by subjectivism in this work is much closer, and perhaps identical, to what Darwall

Sometimes non-cognitivist (or ‘expressivist’) views likewise go under the name of subjectivism. James Rachels, for example, thinks of both the cognitivist/descriptivist view discussed in the previous paragraph (on which evaluative utterances report and describe the speaker’s own attitudes) and the non-cognitivist/expressivist view known as emotivism—on which such utterances are taken instead to *express* the speaker’s attitude—as two different ‘stages’ of subjectivism, with the former view being labelled ‘simple subjectivism’ and taken to make up the first stage, and emotivism is being thought of as a development of that view—a second stage of, and more sophisticated form of, subjectivism.¹⁰ James Ward Smith, in an earlier treatment of the subject, also thinks of these two types of meta-ethical theories as deserving the epithet *subjectivism*.¹¹

means by projectivism. Thus, if Darwall’s interpretation is correct, Hobbes was a subjectivist in my sense of the term. My point about Hobbes is not primarily historical or exegetical. However, since I do take Hobbes to be a major predecessor of value subjectivism in my sense, I will allow myself to make a few interpretative remarks from time to time when this seems to me to be of philosophical relevance.

¹⁰ Rachels (2007). However, when he first introduces the umbrella term of subjectivism, Rachels characterizes it in ontological terms: the ‘basic idea behind Ethical Subjectivism’ is that ‘where morality is concerned, there are no “facts”, and no one is “right”’ (p. 36). But the view he calls simple subjectivism, which is supposed to be one form, or ‘stage’ of subjectivism, clearly is in tension with the ‘basic idea’ thus understood. Indeed, Rachels involves himself in a contradiction when he says, first, that subjectivism implies that no one is ever right, and, then, goes on to criticize one version of subjectivism for having the implication that we are always right (pp. 38–39).

¹¹ Smith (1948). Smith distinguishes three different senses of subjectivism. In his first sense, which he calls ‘traditional’, the label applies to ‘any position which holds that our value judgments are verifiable *always and only* by virtue of the occurrence of people’s mental states such as likings, wants, or satisfactions. Values, it is held, are facts in the world, and our statements may succeed in pointing to them or not’ (p. 393). In the second, ‘positivistic’, sense, subjectivism refers to ‘any position which flatly denies that our value judgments are verifiable in any sense comparable to the way in which our ordinary empirical judgments are verifiable. Values are *not* facts in the world at all; they are fictions we have thrown up in our emotional tantrums’ (p. 394). In Smith’s third sense, subjectivism has nothing to do with verification. Instead, it names the ‘simple conviction that the only thing common throughout valuational situations are the attitudes of evaluators’ (p. 402). Towards the end of his essay, Smith says ‘[T]he single feature common to all of these views is the belief that value theory has for its subject-matter, man. Accordingly, what all such views may be contrasted with is the view that the subject-matter of value theory is *not* man. If man were erased from the universe, the mind of God could still contemplate the subject-matter of value

A subjectivist in my sense (and likewise an objectivist) is not committed to any particular view of the meaning of evaluative utterances, no particular semantic analysis of 'good', 'bad', 'better than', etc. As John L. Mackie says, 'The denial that there are objective values does not commit one to any particular view about what moral statements mean'.¹² The view that I will be discussing under the heading of subjectivism in this work is neutral with respect to cognitivism/descriptivism and non-cognitivism/expressivism. In Chapter 8 I will return to the semantics of evaluative language and discuss which semantic views fit best in conjunction with a subjectivist view about the nature of value. Whilst there need be nothing 'mere' about semantics, my conviction is that the distinct question concerning the subjective or objective nature of value is the central issue in second order value theory¹³ and that this question cannot be settled by an analysis of language.

Saying that the divide between objectivism and subjectivism concerns the nature of value rather than our way of speaking or thinking about it does not straightforwardly reduce it to that between realism and anti-realism about value. Depending on how value realism is understood, there can, at least in principle, be a subjectivist realist as well as objectivist anti-realists. Though, I believe that both of these combinations suffer from internal tensions. The ontological status of value will be discussed in Chapter 7.

One might think that subjectivism cannot be a theory of the nature of value but at best a theory about the nature of valuing. This, I believe, would be a mistake. It might very well be that value subjectivism has implications

theory' (p. 404). Though suggestive, I do not think that this is sufficient to draw a demarcation line between subjectivism and objectivism.

¹² Mackie (1977) (p. 18). Richard Double (2006) makes the point that in order to see that cognitivism/non-cognitivism and subjectivism/objectivism are independent of each other, we need only consider the two conceptual possibilities that one thesis is true and the other false: 'First, subjectivism might be true, even if our best philosophy of language decides that most moral judgments (those other than pure emotive ejaculations) must be treated as indicative statements and, thus, all general noncognitive analyses fail. . . . The second conceptual possibility is that there might exist objective moral facts, even if some variety of noncognitivism fully accounts for the meanings of moral judgments' (p. 5).

¹³ For a different opinion, see Richard M. Hare (1985).

for theories of what it is to value in the sense of thinking or judging valuable (implications that will be further examined later on), but no theses of a linguistic or psychological kind will form part of the core of value subjectivism. This should be clear in the light of the fact that subjectivists and objectivists may hold common views about what it is to think or judge valuable.

Likewise, I believe, may subjectivists and objectivists hold a shared view about the phenomenology of value: of what it is like to value. Subjectivists do not (or, at any rate, need not) claim that how things seem to us provides support for their thesis that value is grounded on attitudes. Indeed, perhaps the most natural view to hold (regardless of the debate between subjectivism and objectivism) is that value has a certain objective-seeming quality. As far as our experience goes, value seems to be an objective feature of the world and not a product of our attitudes.¹⁴ Most subjectivists hold that our experience of value involves an element of illusion. This is, however, denied by some subjectivists. Simon Blackburn writes at one point ‘I would deny that there is an aspect of moral *phenomenology* which gives morality an objective appearance which quasi-realism must regard as illusory’.¹⁵ But even though some may want to say this, subjectivists do not need to say that objectivity is a dispensable aspect of our common experience of value or for that matter that it forms no part of our folk theory of the nature of value. Subjectivists and objectivists alike can, for example, hold the view that a claim to objectivity forms an essential part of our evaluative thought and language. The point I wish to make here is merely that since subjectivists and objectivists can (though need not necessarily) hold a common view about the phenomenology of value, this cannot be the locus of the disagreement between them. The debate over the phenomenology of valuing, like that of the semantics of evaluative language, is separate from that of the nature of value and it is the latter debate that subjectivists and objectivists are engaging in.

¹⁴ For a contrary view, see Loeb (2010).

¹⁵ Blackburn (1985) (p. 21, n. 8). Though, Blackburn also says things that seem to point in the opposite direction. For example, Blackburn concedes (to Mackie) that it ‘may be true’ about ‘most ordinary moralists’ that they ‘falsely imagine a kind of objectivity for values, obligations, and so on’ (ibid, p. 3).

IV. Supervenience

Subjectivists and objectivists agree that value is something that accrues to objects¹⁶, and that it somehow depends, or ‘supervenes’, on the ordinary, natural, or non-evaluative, features of these objects. It is widely thought that whatever value a particular object is taken to possess, it must be seen to have this value in virtue of some of its other properties (other than its goodness or badness, etc.). If a particular painting is considered beautiful, for instance, it must be held to be so in virtue of having certain features (other than beauty itself), such as its particular colours, shapes, and proportions, its originality and authenticity, and so on. In other words, if the painting is beautiful, there must be something about the painting that makes it so; there must be something about the painting that the person evaluating it can point to when asked ‘Why is this painting beautiful?’ In general terms, the thought is this: necessarily, if an object is (thought to be) valuable, this is so in virtue of some of its non-evaluative properties.¹⁷ These properties (whatever they may be) must be taken to make the object in question valuable. It is common to call these properties the value-making properties of the object and to say that together they make up the supervenience base of the value (though this will be problematized shortly, as there may be more to supervenience than value-making features).

This much could be taken to be common ground among subjectivists and objectivists, since it seems to form part of what they are both trying to explain; it forms part of that shared analysandum of which subjectivists and

¹⁶ The term *object* is here used in a way that differs somewhat from how it is used in everyday language. Anything that can carry value is an object in the relevant sense. The issue of what kinds of objects can be bearers of value will be discussed below, but will ultimately be left open.

¹⁷ As Dan Egonsson has suggested to me, saying that value always supervenes on non-evaluative features of valuable objects is not entirely satisfactory because it seems to exclude the possibility that some evaluative features may supervene on other evaluative features. This is a possibility I want to leave open. Perhaps there can be several layers of evaluative features with each layer depending/supervening on the layer just below it. However, I believe that at the bottom of such a layered structure we will always find some non-evaluative features upon which the first evaluative layer depends. The layered structure cannot be evaluative all the way down.

objectivists propose rival analyses. In light of the ambition to identify a common ground between subjectivists and objectivists, it is therefore tempting to insist that both subjectivists and objectivists can talk about value supervenience (and mean the same thing). Though, since they have different views about the nature of value, they naturally will have different views about the precise nature of this relation. Insofar as this is so, supervenience could be used as a shared ‘litmus test’ of plausibility for theories of value. Any plausible such theory must be able to account for supervenience, and the better the account the better the theory in that respect. This strategy has been employed by some subjectivists, arguing that a subjectivist theory provides a better explanation of supervenience than objectivist rivals.¹⁸

The feature of supervenience that is of particular importance for my purposes is the following consistency requirement. If an object O_1 is valuable and a different object O_2 is not valuable, then this must be because of some difference in the non-evaluative properties of O_1 and O_2 . If it is agreed that there is no non-evaluative difference between O_1 and O_2 , it must also be agreed that there is no evaluative difference between them. This idea goes back at least to G. E. Moore, who wrote that ‘[S]upposing a given thing possesses [intrinsic value] in a certain degree, anything exactly like that thing must possess it in exactly the same degree’.¹⁹ This consistency requirement is rooted in our pre-theoretical intuitions. When a person is asked why they value a certain object but not another, we normally expect them to be able to point to some difference between the two objects other than that one is good and the other bad (or whatever values are being ascribed to the two objects). I will assume that subjectivists and objectivists alike want their theories to respect this requirement.

It might be thought that subjectivists, who propose to analyse value in terms of attitudes, would have a problem with this consistency requirement. One might insist that it is quite possible for an evaluating subject to have an attitude towards O_1 (in virtue of some of its features), but lack a similar attitude towards O_2 (in virtue of the same features), even though the two objects are exactly alike in all their non-evaluative features. Wouldn’t

¹⁸ See, for example, Mackie (1977) (p. 41) and Blackburn (1984) (pp. 182–187). I will return to their arguments in a later chapter.

¹⁹ Moore (1922) (p. 269)

subjectivism then lead to a breakdown of consistency? I take up this challenge in Chapter 3 and argue that it does not need to do so.

Enablers and Disablers

Some philosophers have argued that there is more to supervenience than value-making features.²⁰ Jonathan Dancy stresses that there are several different ways in which the presence or absence of a feature can make a difference to the value of something, Dancy's main distinction being between favouring/disfavouring and enabling/disabling something else to favour.²¹ Dancy introduces the term *resultance base* to cover those properties of the object that make it good (that is, the value-making features, the favourers). This is different from the *supervenience base*, which includes all the non-evaluative features of the object, not just those that make it good. Hence, the supervenience base is something far more extensive than the resultance base. The supervenience base 'contains also every other feature (of the value bearer, but perhaps also of other things) that can in any circumstances make a difference to the ability of the features in the resultance base . . . to play their special [value-making] role'.²²

Dancy's notions are not going to play a large role in my discussion, though I will have reason to return to them occasionally. As far as I can see, these notions, and the arguments for recognizing them, can be adopted by subjectivists as well as by objectivists. Dancy briefly considers the idea that the motivation behind what he calls 'atomism'—the view that a feature that is a value-maker in one case must remain a value-maker in any other case—is 'an adherence to the view that reasons for action are partly grounded in desires'. This is so, Dancy suggests, because 'if we accept that view, and if we then think of desires as giving the desirer the same reason wherever the desire occurs, the result looks atomistic'. I am not sure, however, why an adherent of such a view would necessarily come to that conclusion. Dancy

²⁰ For example, Dancy (2004).

²¹ Ibid (pp. 38-39). Dancy also speaks of considerations playing the roles of intensifiers and attenuators (pp. 41-42).

²² Ibid (p. 178)

himself points out that it could be responded to this that ‘even if all practical reasons are grounded in desires, the same desire need not always function as the same reason’.²³ What it means for reasons and values to be grounded on desires is less clear, and I will return to that claim below. Here I just wish to say that in the way that I understand that claim, it does not seem to constitute an obstacle to Dancy’s holism—the view that a feature that is a value-maker in one case need not be a value-maker, or even be a disvalue-maker, in another depending on the presence or absence of other features that enable or disable these features to be value-makers.

V. Final and Intrinsic Value

That something is *finally* valuable I will understand as meaning it is valuable for its own sake. Arguably going back at least to Aristotle, final value is often contrasted with instrumental value—the value of something as a means to something else.²⁴ That something is *intrinsically* valuable is usually understood as it being valuable in itself. In the tradition of G. E. Moore, intrinsic values ‘depend merely on the intrinsic nature of [the valuable object]’.²⁵ In the language of supervenience, an object’s intrinsic value supervenes exclusively on the object’s internal properties.

It has traditionally been assumed that the final value of something must be intrinsic. That is, insofar as some object is value for its own sake it must possess this value in itself, in virtue of its intrinsic nature; the final value of an object must supervene on its internal properties. During the last few decades, however, several philosophers have argued that final values

²³ Ibid (p. 75)

²⁴ However, there might be other kinds of non-final values in addition to instrumental ones (such as contributive values).

²⁵ Moore (1922) (p. 256)

need not be intrinsic.²⁶ On this approach, the final value of an object can supervene not only on its internal properties, but also on its relational properties. Two guitars, for example, may have exactly the same internal properties but still differ with respect to their final values. This could be so because only one of the guitars used to belong to Steve Howe. If the relation to Howe is what makes the guitar valuable for its own sake, this would be an example of a final value that is not intrinsic—a final extrinsic value.²⁷

It has been thought by some philosophers that only objectivists can talk about intrinsic value. Moore wrote '[F]rom the proposition that a particular kind of value is "intrinsic" it does follow that it must be "objective"'. And he further thought 'It is this conviction—the conviction that goodness and beauty are intrinsic kinds of value, which is, I think, the strongest ground of [the] objection to any subjective view'.²⁸ As Wlodek Rabinowicz and Toni Rønnow-Rasmussen argue, this is a mistake, one that they point out Shelley Kagan and Christine Korsgaard also make.²⁹ Kagan and Korsgaard both agree that there are both extrinsic and intrinsic final values, but they argue that for a subjectivist there can be no intrinsic values. Kagan writes, for example, that on 'a radically subjectivist conception of value, according to which nothing would be valuable as an end in the absence of there being some creature who values it' nothing can be

²⁶ See, for example, Korsgaard (1983), Kagan (1992), O'Neil (1992) and Rabinowicz & Rønnow-Rasmussen (2000). Rabinowicz and Rønnow-Rasmussen trace the view back to Beardsley (1965). For a critical discussion, see Chapter 3 of Zimmerman (2001).

²⁷ My example of Steve Howe's guitar is closely modelled on examples suggested by Rabinowicz & Rønnow-Rasmussen (2000). They suggest that Princess Diana's dress, Napoleon's hat and a gun that was used at the battle of Verdun, are possible candidates for carrying final extrinsic value (p. 41). Any of these objects might be valued for its own sake in virtue of its special relationship to a particular person or event. Other examples of final extrinsic values from the literature include objects that are valuable for their own sake on account of relational properties such as their uniqueness, rarity or originality. A rare stamp may, for example, be valued for its own sake precisely in virtue of being rare. A work of art may similarly be valued on account of its originality. John O'Neill (1992) gives as an example that a wilderness may be valuable on account of being untouched by human hands. Whether any of these objects are valuable at all is, of course, a first order question that need not concern us at this point.

²⁸ Moore (1922) (p. 255). It may be noted, however, that Moore did not understand subjective and objective in the same way that I understand them here.

²⁹ Rabinowicz & Rønnow-Rasmussen (2000) (pp. 39–40)

intrinsically valuable.³⁰ Korsgaard, in a later piece, similarly writes that ‘[our particular ends] have only extrinsic value, since their value depends on our own desires and interests in them and is conferred on them by our own rational choices’.³¹ The thought is a natural one. After all, if the subjectivist holds that the value of an object is always conferred upon it by some attitude that some subject has towards this object, how can it at the same time be held that the object in question has its value solely in virtue of its internal features? If the values of objects always hinge on something external to the objects themselves (on the attitudes that are being directed at them), how can there be intrinsic values? The short answer is that it is only the features in virtue of which the object is desired (or otherwise favoured or disfavoured³²) that need be internal to the object in order for the conferred value to be intrinsic. The long answer will have to wait until the next chapter, after some further terms and distinctions have been introduced.

Because I take final values (whether intrinsic or extrinsic) to be the normatively interesting type of value (in the sense of involving reasons for certain responses; see Section VII below), I will be primarily concerned with final values. But a subjectivist analysis of intrinsic value along the lines hinted at above will be outlined in Chapter 2. Understanding the notion of

³⁰ Kagan (1992) (p. 184). In his (1998) Kagan reiterates the point. He notes first that ‘the subjectivist does believe that many objects do indeed possess value as an end’, or for its own sake; what I am here calling final value. But he then goes on to say: ‘Obviously, however, they do not possess that value solely by virtue of their intrinsic properties. For it is not an intrinsic property of an object that it is valued by some creature’ (p. 281). Kagan thus thinks that insofar as we take intrinsic value to mean value that something has solely in virtue of its intrinsic properties, a subjectivist must hold that nothing at all has intrinsic value. Kagan is, however, reluctant to accept the conclusion that subjectivists cannot recognize intrinsic values and proposes instead that we allow for talk of intrinsic value under subjectivism by giving up the idea that intrinsic values must depend solely on the intrinsic properties of their objects (ibid). The conclusion Kagan wants us to draw is that ‘according to subjectivism, intrinsic value depends upon a relational property’ (p. 282). As indicated, I think that giving up on the traditional (Moorean) understanding of intrinsic value as depending solely on the intrinsic features of the value bearer would be a mistake. What Kagan fails to appreciate is that subjectivists can recognize intrinsic values in the traditional sense. There is thus no need to give up, or rethink, the traditional understanding in order to allow for talk of intrinsic value under subjectivism. See also Rabinowicz & Rønnow-Rasmussen (2000) (p. 49, n. 29).

³¹ Korsgaard (1998) (p. 63)

³² I will use favouring and disfavouring as placeholders for a wider range of different pro and con attitudes, including but not limited to desire and preference.

intrinsic value in a particular way is, as far as I can see, not an essential commitment of either subjectivism or objectivism. But holding on to the traditional, Moorean understanding of the notion, whilst showing that subjectivists can analyse intrinsic value in that sense, will straightforwardly allow us to maintain the guiding assumption that subjectivists and objectivists are offering rival analyses of the same thing. Both subjectivists and objectivists can legitimately speak about intrinsic value, and mean the same thing by that: a value that depends/supervenes exclusively on the internal properties of its bearer. Of course, beyond this common understanding of the notion of intrinsic value, subjectivists and objectivists offer fundamentally different analyses of its nature. One implication of subjectivism is that just as an object may be valuable relative to one person but not relative to another, so too may one and the same object have different kinds of value relative to different people. The guitar that once belonged to Steve Howe may, for example, be finally valuable relative to one person, but merely instrumentally valuable relative to another; or it may be intrinsically valuable relative to one person whilst extrinsically valuable relative to another—depending, as we shall see, on their respective attitudes. That value may thus accrue to an object relative to one person but not relative to another is an aspect of the subjectivist analysis that may be quite alien to objectivists.

VI. Value Bearers

If the guitar that once belonged to Steve Howe is taken to be valuable for its own sake, then concrete objects must be counted among the possible bearers of final value (since guitars are concrete objects). Some philosophers, however, have held that only states of affairs can carry final value.³³ This issue is connected to the one that was discussed in the previous section about whether all final values are intrinsic. If only states of affairs can carry final value, it is natural to think that all final values are intrinsic: that they always

³³ Michael Zimmerman is a case in point. See his (2001) (pp. 33–39).

depend/supervene on the internal properties of the state in question. When I say that Steve Howe's guitar is valuable for its own sake in virtue of once having belonged to Howe, it could be insisted that what I'm really getting at is that the state of affairs *that* the guitar has belonged to Howe or the fact that the guitar that has belonged to Steve Howe *exists* is what carries final value.³⁴ As this state contains the object as well as its relation to Howe as component parts, the value of the state depends solely on the internal properties of the state itself and as such is an intrinsic value. If claims to the effect that something possesses final value in virtue of its relational properties were always reducible to claims that some intrinsically valuable state of affairs involves the object in question as well as its relations to other things, there would be no need to postulate extrinsic final values.

But as Rabinowicz and Rønnow-Rasmussen have argued, such reductions are not reasonable. First, as these authors point out, those attempting the reduction would need to tell us which state it is that is supposed to carry the value. Is it the state of affairs that the object exists? Or that it continues to exist given that it already exists? Or that people preserve it? There seem to be many possible candidates and it is hard to see how the right one could be singled out. Second, the reduction manoeuvre seems to put the cart before the horse: if the existence of a certain object is valuable, it is because the object itself is valuable, not the other way around. The value of the object is scarcely explicable in terms of the value of the state.³⁵ Third, it seems quite possible to hold that a certain object is valuable while at the same time deny that the existence of this object is a valuable state.³⁶ Fourth, and finally, Rabinowicz and Rønnow-Rasmussen argued at length that the reduction manoeuvre is insufficiently motivated in the first place.³⁷ Why would we even want to reduce the value of concrete objects to abstract states of affairs? With respect to the first of these four points, a value-bearer monist

³⁴ Zimmerman (ibid) argues that the attribution of value to concrete objects is 'eliminable in terms of, or at least reducible to, an attribution of final value to some states of the objects' (p. 39).

³⁵ Rabinowicz & Rønnow-Rasmussen (2000) (pp. 43–44)

³⁶ As Rabinowicz and Rønnow-Rasmussen (ibid) point out, in *Kritik der Urteilskraft* (1790) Immanuel Kant held that the judgment that something is of aesthetic value does not in the least commit one to say that the object's existence is of value (§2).

³⁷ Rabinowicz & Rønnow-Rasmussen (ibid) (pp. 44–49)

who is also a subjectivist has a reply. The state that is the value bearer in a particular case is the state that is the object of the relevant value-conferring attitude. If the object of the relevant attitude is the state of affairs that something exists, for instance, then it is this state that carries the value in question; if it is instead some other state that is the object of the attitude, then that state is the value bearer. This response may, however, backfire if it turns out that the objects of the relevant attitudes include not only states of affairs but also concrete entities, persons, etc.

Value-Bearer Monism and Value-Bearer Pluralism

In general terms, those who hold that only objects belonging to a single ontological category can carry final value are value-bearer monists, whilst those who allow that value may accrue to objects belonging to several different ontological categories are value-bearer pluralists. The distinction between value-bearer pluralism and value-bearer monism should not be conflated with the distinction mentioned in Section I between axiological pluralism and axiological monism. Axiological pluralists endorse a plurality of values—they hold that several different things, such as knowledge, peace, biodiversity, friendship, pleasure or music, are finally good and bad, better and worse, etc. Value-bearer pluralists, by contrast, endorse a plurality of possible bearers of value. Candidates include abstract entities such as states of affairs, facts, and (instantiations of) properties, concrete objects such as persons and paintings, and events such as experiences.³⁸

An axiological pluralist need not be a value-bearer pluralist. One can, for instance, hold that there are several distinct final values that all accrue to states of affairs. If knowledge, peace, and biodiversity are held to be valuable for their own sake, for example, it would always be the states of affairs *that* things are known, *that* peace is kept, and *that* biodiversity is preserved, that would carry these values. Perhaps neither need an axiological monist be a value-bearer monist? A hedonist, someone who thinks that pleasure alone is valuable for its own sake, might hold that the bearers of final value are not restricted to experiences but may include also states of affairs in which pleasurable experiences take place.

³⁸ I thank Frits Gävertsson for helping me with this list.

Subjectivists and objectivists can be either value-bearer monists or value-bearer pluralists. Insofar as subjectivism is formulated exclusively in terms of preferences, it is perhaps natural to think of the objects of value as being restricted to states of affairs, and thus to think of the subjectivist as adhering to a version of value-bearer monism.³⁹ However, we often say things like ‘I prefer an apple to a pear’ which sounds like a preference for one concrete object over another. This particular preference can perhaps be understood as a preference for the state of affairs in which I eat an apple over the state in which I eat a pear, but it is not obvious that a reduction of this sort is always plausible.

A more secure route to value-bearer pluralism opens itself up for subjectivists as soon as they allow a wider set of pro and con attitudes to ground values. In addition to desire and preference, they may want to include, for example, admiration, adoration, cherishing, honouring, respect, love, etc. Insofar as subjectivists are willing to operate with such a broader set of attitudes, then they should also be willing to allow that value may accrue to objects belonging to several different ontological categories. After all, it sounds odd to admire or to respect a state of affairs, for example.⁴⁰ An argument of this kind is available to anyone (subjectivist or objectivist) who thinks that value should be analysed in terms of the responses taken (or that it is appropriate to take) towards the valuable object.⁴¹

³⁹ Gauthier (1986) is a case in point. He doesn’t explicitly endorse value-bearer monism, but as will be seen later he characterizes his subjectivist value theory in terms of preferences for states of affairs.

⁴⁰ Though, Zimmerman (2001) thinks that this may be due just to an accident of language (p. 43).

⁴¹ Rabinowicz & Rønnow-Rasmussen (2000) admit that if to be valuable is to be ‘a *fitting object of a preference* (or, in a more subjectivist version of such a preferentialist position, it is to be an object of a preference we would have under certain ideal conditions)’ or, if value is analysed in terms of pursuing or promoting; if ‘value is *what we should promote* (or what we would promote under ideal circumstances)’, then it would be plausible that states of affairs are the only bearers of value (p. 45). But, as they go on to argue, ‘if we accept the general idea that value is what calls for an appropriate response, the response in question need not consist in just preferring or promoting’ (p. 46). If we include love, liking, admiration, respect or behavioural responses such as cherishing, protecting and caring, then ‘the main motivation to reduce thing values to state values disappears’ (ibid).

Is There a Presumption in Favour of Value-Bearer Pluralism?

It has been argued that if the value-neutral character of second order value theory (which I have emphasised above) is to be preserved, then value-bearer monism must be rejected. That is, if the commitment to neutrality among first order views is to be respected on the second order level, then we ought to endorse value-bearer pluralism. Toni Rønnow-Rasmussen, for example, has argued that if it would be held that only states of affairs can carry value, certain first order axiological views would seem to be excluded from our possibilities.⁴² Take the example again of axiological hedonism, a form of axiological monism that ascribes value exclusively to pleasure. Suppose that this view ascribes final value exclusively to experiences and not to states of affairs at all. Such a first order axiological view would be ruled out by a value-bearer monism that allows only states of affairs to carry value. Whatever other arguments might convince us to reject this kind of hedonism, it does not seem right to rule it out on the second order level. On the other hand, if it were held that only experiences can be bearers of value, many other axiological views would thereby be illegitimately excluded. Either form of value-bearer monism would involve a breach of the value neutrality of second order views. The upshot is that if second order value theory is to leave the choice between different first order views genuinely open, then value-bearer monism cannot be maintained. This argument, if accepted, seems to establish a presumption in favour of pluralism with respect to value bearers that any would-be monist would need an argument to overcome.

A potential concern about this argument is that adopting value-bearer pluralism in a similar way rules out monistic axiological views. Consider once more the case of the hedonist and assume that all objects being ascribed value by this view belong to the same ontological category (experiences, presumably). Such a hedonist, so the concern goes, cannot be a value-bearer pluralist. Holding that the only valuable objects are pleasurable experiences would mean ascribing value to objects in one ontological category only and thus involves rejecting value-bearer pluralism after all. Since this is so, the concern continues, remaining neutral with respect to

⁴² Rønnow-Rasmussen (2002)

monistic and pluralistic axiological views is not compatible with endorsing pluralism with respect to value bearers. The only way to preserve first order neutrality would be to also remain neutral with respect to value bearers.

I believe that this issue hinges on how value-bearer monism and pluralism are more precisely to be understood. If we interpret these views as strictly second order views without any first order commitments, then a monist is someone who holds that there are formal (conceptual and/or metaphysical) obstacles to ascribing value to objects from more than one ontological category. Meanwhile, a pluralist is anyone who denies the existence of such formal obstacles and thus allows that, as far as second order considerations go, it is possible for value to accrue to objects belonging to more than one ontological category. Interpreted in this way, none of these views forces us to value anything in particular, and value-bearer pluralism thus understood does not force us to value objects belonging to at least two different ontological categories. Given this interpretation, such a pluralist can (contrary to the assumption underlying the concern above) be an axiological monist: a hedonist, for example. This would involve holding that, as far as second order value theory is concerned, objects belonging to several different ontological categories may be bearers of value; there are no formal obstacles to ascribing value to several different types of objects. But, moving to the first order level, and thus taking a certain normative or evaluative stance, this hedonist can claim that only one kind of object—pleasurable experiences—actually has value. A hedonist who was such a value-bearer pluralist but an axiological monist would say that someone who ascribes value to objects that are not experiences is substantially mistaken on the evaluative level but not conceptually or metaphysically confused.

A proponent of value-bearer monism understood in this formal way could not, as far as I can see, credibly make the corresponding claim that, since it too is a second order view, it too leaves first order questions open. Saying that it is impossible for objects belonging to more than one ontological category to carry value does preclude us from valuing objects that fall outside the category in question, and that does seem to rule out certain first order axiologies. It is upon this asymmetry that the aforementioned presumption in favour of value-bearer pluralism rests. But this presumption only favours pluralism understood in such a formal way. There is no presumption in favour of pluralism interpreted as a first order view, or as entailing certain first order commitments, thus forcing us to ascribe value to objects belonging to at least two different ontological categories. At least, no such presumption has been established. In the end,

the only way to preserve value neutrality would then be to remain neutral with regard to the nature of value bearers.

VII. Value and Normative Practical Reasons

Practical reasons are often contrasted with theoretical reasons, such as reasons for belief. A practical reason is a reason for doing something, for acting in some way rather than in another.⁴³ What practical reasons we have, what we have reason to do, or how we should or should not act are examples of first order, normative questions and consequently fall outside the scope of the present study. It is commonly held that practical reasons are considerations that ‘count in favour’ of some action. That Ulan Bator is the capital of Mongolia can be a practical reason for someone if it is a consideration counting in favour of some action (say, going there for someone who has the personal goal of visiting every capital in Asia, or counting in favour of not going there for someone who is wanted by the Mongolian authorities).

Some philosophers think that it is impossible to provide an analysis of the concept of a reason and the meaning of the expressions ‘a reason’, and ‘being a reason for’. T. M. Scanlon, for example, takes the idea of a reason as primitive and thinks that the idea of a reason ‘does not seem to . . . be a problematic one that stands in need of explanation’. ‘Any attempt to explain what it is to be a reason for something’, Scanlon claims, ‘seems to me to lead back to the same idea: a consideration that counts in favour of it’.⁴⁴ Derek Parfit similarly thinks that the concept of a reason is indefinable and thus cannot be helpfully explained ‘merely by using words’ (though he too holds that ‘gives a reason for’ means ‘roughly the same’ as ‘counts in favour of’).

⁴³ We can understand *action* here in a broad sense to include not only overt actions (opening the door, picking up a book, kicking the ball, etc.), but also mental acts such as choosing and trying. For my present purposes we could also count reasons for having certain attitudes as practical reasons.

⁴⁴ Scanlon (1998) (p. 17)

Parfit further says that the concept of a reason is fundamental like those involved in our thoughts about time, consciousness and possibility.⁴⁵

As I did in Section III above, I think it is vital that we distinguish between semantic and ontological analysis. To repeat: subjectivists and objectivists are not committed to any particular semantic analysis of ‘good’, ‘bad’, ‘better than’, etc. Similarly, they are not committed to any particular semantic analysis of ‘a reason’. But this must be kept apart from the fundamentally different project of accounting for the grounds on which a fact attains the normative status of being a reason. The latter project is, in the words of Jonathan Dancy, about ‘the metaphysical ground for a reason—about what creates or generates it’.⁴⁶ On what grounds can a fact or true proposition⁴⁷ be said to count in favour of some action? This is something that needs to be explained and this is a point on which I take subjectivists and objectivists to diverge.

We may think of this latter project concerning reasons as analogous to the one about value that I have been describing in previous sections. Just as a non-evaluative property of an object can be value-making, so too can a natural fact be reason-giving and thus count in favour of some action. This much I take to be common ground among subjectivists and objectivists. But the normative status of being a value-maker or a reason-giver is something that needs to be explained, and it is here that subjectivists and objectivists differ. Very roughly, subjectivists deny—and objectivists affirm—that normative status, like value, is an objective phenomenon. All agree that it is straightforwardly a fact that Ulan Bator is the capital of Mongolia, but can it similarly be an objective fact that this fact is a normative reason, that it counts in favour of something? As I will argue later on, subjectivists hold that things can be good or bad, or have the normative status of being a reason, only relative to subjects. Objectively, nothing is good or bad, and nothing counts in favour of anything. It can only do so from the perspectives of valuing subjects.

⁴⁵ Parfit (2011) (p. 31)

⁴⁶ Dancy (2000) (p. 17)

⁴⁷ I will not take a stand on whether reasons are facts or true propositions, but I will drop the ‘or true propositions’ in what follows.

A Proposal

I turn now to the central issue of this section: what is the relation between reasons and values? I suggest that subjectivists and objectivists should agree on the following conditional:

Whenever some object *O* is valuable, there are reasons to respond to *O* in some way or other, where those reasons derive from the features of the object in virtue of which it has its value.⁴⁸

It is further thought that insofar as the value that *O* carries is a positive value (for example, goodness), the reasons in question are for responding to *O* in some positive manner, and if *O* is of negative value (for example, badness), some negative response is called for. It is commonly held that something's being good (for example) is not itself a reason to respond positively to it; rather, it is the features that make the object good that count in favour of a positive response. I take the formulation I have proposed to be neutral between the view that these reasons are value-makers and the view that these reasons are the values themselves (the latter also deriving from value-makers).

Two further elements of the formula need to be immediately clarified: the expressions *respond to* and *there are reasons*. The latter expression might seem to imply that reasons are some kind of free-floating normative entities that need not relate to agents. However, it is plausible to assume that a reason is always a reason *for* someone or other. The formula does not specify *for whom* there are thought to be reasons, since this might vary from case to case depending on the type of value and the type of object involved, among other variables. Some reasons might be reasons for everyone, or at least for anyone who is capable of recognizing the reason, perhaps restricted to those who happen to be properly situated, have the ability to perform the action that is being called for, etc. Other reasons might be reasons only for a select few who perhaps stand in some special relation to the valuable object in

⁴⁸ Though the conditional proposed above may not strictly speaking be common ground among everyone involved in the debate, I think it is a fruitful starting point. A potential example of a detractor from my linking of the evaluative and the normative is W. D. Ross (1930) who argued that 'the right' and 'the good' are quite separate.

question.⁴⁹ Here there is room for a spectrum of different views that can be accepted by subjectivists and objectivists alike. The disagreement between subjectivists and objectivists, as it pertains to reasons, concerns the subjective or objective status of normative reasons, and this should not be conflated with the issue regarding to whom reasons apply, i.e., who has the reasons. Subjectivists claim (which will be elaborated on later) that the status of being a normative reason is always relative to some subject's perspective; it is important to recognize that this claim is silent on the issue of *for whom* it is a reason. For instance, relative to a particular subject's perspective, everyone might have a certain reason. Though the above formula can be accepted by subjectivists and objectivists alike, they interpret it in different ways. In a subjectivist reading, it is always value relative to subjects that implies reasons relative to subjects.

The second thing I wanted to comment on concerns the expression *respond to*. This expression is deliberately vague in order to encompass various different views. Some think that the way in which we have reason to respond to a valuable object is by having some attitude towards the object in question. If some person is admirable, for example, then there is reason to admire him; if something is desirable, then we have reason to desire it, and so on. It might be thought that this interpretation of *respond* would fit uncomfortably with subjectivism because the combination might seem to lead to the following problem. If value is determined by attitudes, it would seem that whenever someone has an attitude, he or she therefore also has a reason to have this very attitude; one always has reasons to have the attitudes that one does in fact have!

However, I believe that this is less troublesome for subjectivists than it at first may seem (though it might be a serious problem for some non-subjectivist views that understand reasons in terms of attitudes). Let me explain. First, as will become clearer later on, it is neither the attitude nor

⁴⁹ Sometimes the former are called *agent-relative* and the latter *agent-neutral* reasons. These terms are, however, used also in other ways. Thomas Nagel (1986) proposed that an agent-neutral reason is a reason that 'can be given a general form which does not include an essential reference to the person who has it'. If, on the other hand, 'the general form of a reason does include an essential reference to the person who has it' it is an agent-relative reason (pp. 152–153). Subjectivists and objectivists may agree that there are both kinds of reason in Nagel's sense. I will use the term *relative* in yet a different way (contrasting it with *absolute*).

the fact that I have a certain attitude, nor (necessarily) the goodness itself (that the attitude is thought to bestow on the object) that justifies or makes the attitude fitting. Rather, it is the non-evaluative features of the object *for which* I favour it that, in the light of my attitude, are seen as reason-giving. It is to these features of the object, rather than to the attitude, that the subject will naturally point when asked for his reason for valuing the object in question. When we wonder whether we are justified in having a certain attitude—whether we have reason to have the attitude in question—the fact that we do now have this attitude is moot. It is instead something about the object of our attitude that is seen to merit the attitudinal response.

Second, on subjectivist views, as I have insisted and will argue for later on, things are valuable only relative to perspectives, and similarly with reasons: some fact is a normative reason only relative to a perspective. Hence, if we interpret *respond* as having an attitude, when something is valuable relative to a particular perspective, whilst it is the case that—from this perspective—there is a reason to favour this object, subjectivism does not imply that by having an attitude you thereby make it the case that you have a reason to have this attitude from any *other* perspective. Thus, on subjectivist views as I understand them here, it is not the case that we can derive the evaluative conclusion that one has a normative reason to have a certain attitude from the fact that one now has this attitude. We can only derive the non-evaluative conclusion that one has such a reason relative to one's perspective. I therefore think that interpreting *respond* as having an attitude does not lead to an obviously objectionable type of evaluative 'bootstrapping' within a subjectivist framework.

If you believe that—in the same way that *ought* implies *can*—*having a reason* implies *can*, and you also believe that what attitudes you have is not 'up to you' in the relevant sense, then I take it that you will find it unattractive to interpret *respond* as having an attitude (independently of the bootstrapping concern). An alternative way of thinking about responding to a valuable object would be in terms of taking some action towards it. Candidate actions include choosing, procuring, prolonging, preserving and so on (for objects of positive value), and abolishing, avoiding, destroying, preventing and so on (for negative values). Which action one would have reason to take on a particular occasion depends on the type of value and the type of object involved. Insofar as we have reasons to protect or destroy something, that something is plausibly a concrete object like a building or an artefact, not an abstract entity like a state of affairs. And insofar as we have reasons to prolong or prevent something, that something is plausibly an

event like an experience or a war, and not a concrete object.⁵⁰ For those who understand *respond* as taking action, a practical normative reason would then always be a reason for acting. I will not take a stand on which of the two interpretations is preferable.

Note on the Wrong Kinds of Reasons Problem

Another thing should be noted about the suggested relation between values and reasons. The formula I offered specified a conditional: if something is valuable, then there are reasons. This much—and *only* this much—I take to be common ground among subjectivists and objectivists. Some philosophers hold the stronger thesis that there is a conditional running in the other direction as well: that is, whenever there are reasons there is value. If this second conditional holds, then whenever there is a reason to respond to an object in some way, this object is valuable. A potential problem with accepting this is that it runs into the so-called ‘wrong kinds of reasons problem’.

This problem has been discussed specifically as it pertains to the fitting-attitude view of value. On the fitting-attitude view, if we have reason to have a final⁵¹ pro attitude towards a certain object, then that object is finally good.⁵² To see how the wrong kind of reason problem arises for this view, imagine a situation where you have reason to have a final positive attitude towards a paradigmatically worthless thing. Envisage, for instance, an evil demon who demands that you admire him for his own sake or else suffer severe punishment. In such a case, it seems that you have an excellent reason to admire the demon for its own sake. Given the fitting-attitude view, this leads to the conclusion that the demon is admirable. But such a despicable creature is clearly not admirable. Something has gone wrong: the reason you have to admire the demon (to avoid severe punishment)—powerful though this reason is—does not seem to be the right kind of reason

⁵⁰ Insofar as we have reasons for actions of several kinds towards objects that are valuable, this might perhaps give some further support to value-bearer pluralism.

⁵¹ By ‘final pro attitude’ I mean an attitude of favouring something for its own sake.

⁵² Indeed, on this view—which has been proposed as an analysis of final goodness—this is what it is to be finally good.

to ground admirableness. The fitting-attitude view seems to lack the ability to differentiate wrong kinds of reasons from right kinds of reasons.⁵³

A defender of the fitting-attitude view might avoid this problem by denying that, in cases such as the one described, we really have reason to admire the demon. Perhaps we have reason only to *try* or to *want* to do so. Here the choice between understanding *respond* as having an attitude versus as taking action becomes relevant. If we take the relevant responses to be actions rather than attitudes, the problem cannot be avoided in this way. It cannot so easily be denied that one would have good reason to protect and defend the demon or to genuflect before him in order to avoid severe punishment, even though he is a worthless and despicable creature.

The literature offers several other solutions to the wrong kind of reason problem. As noted, the conditional formula I propose in the previous subsection does not involve us in this problem. It remains, however, a potential problem for those subjectivists and objectivists who accept the biconditional relation between reasons and values.

VIII. Good and Good For

Sometimes the expression ‘good for’ is used in a descriptive as opposed to evaluative or normative way. ‘The heart is good for pumping blood’ can be interpreted as a purely non-evaluative statement. Such an interpretation merely says that the heart plays a certain functional role in the body. ‘A good heart’ (an example of an attributive use of ‘good’⁵⁴) is presumably a heart

⁵³ The example of the demon is from Rabinowicz & Rønnow-Rasmussen (2004), who attribute an earlier version of it to Roger Crisp. Crisp (2000) writes: ‘Since there can be reasons for preferring things that have nothing to do with their value, definitions in terms of preferability are false. Imagine that an evil demon will inflict a severe pain on me unless I prefer this saucer of mud; that makes the saucer well worth preferring. But it would not be plausible to claim that the saucer of mud’s existence is, in itself, valuable’ (p. 459).

⁵⁴ Peter Geach (1956) distinguished between attributive and predicative uses of ‘good’. The attributive use can be elucidated by saying that ‘good’ is like ‘big’ in that the expression ‘X is a good G’ is not equivalent to ‘X is good’ and ‘X is a G’ just like ‘X is a big G’ is not equivalent to ‘X is big’ and ‘X is a G’. Consider Geach’s own example ‘X is a big flea’. This statement is

that performs this function efficiently. What about ‘exercise is good for people’? Insofar as this statement simply means that exercise promotes the effective functioning or longevity of the human organism or some such thing, then the statement does not seem to be an evaluative or normative one.⁵⁵ It seems as if we can say such things without necessarily implying that there are normative reasons to respond in any particular way.

Consider the following example by Adam D. Moore:

We might consider the claim that nitrogen, in certain quantities, is objectively valuable for many varieties of plants. Such plants have a specific nature, and there are certain states of the world that will sustain or promote

not equivalent to ‘X is a flea’ and ‘X is big’. By contrast, when ‘good’ is used predicatively, it functions instead like ‘red’ in ‘X is a red book’. This statement can readily be split up into ‘X is a book’ and ‘X is red’ (p. 33). Geach’s examples of predicative uses of ‘good’ (and ‘bad’) include ‘pleasure is good’ and ‘preferring inclination to duty is bad’ (p. 35). Geach argued that just as there is no such thing as being simply big, but rather only being a big such-and-such, there is no such thing as being plainly good, but rather only being a good such-and-such. Thus, some ‘substantive has to be understood’. Cf. Judith Jarvis Thomson’s (1997) treatment. Geach holds further that ‘good’ and ‘bad’ are always used attributively and never predicatively. Interestingly, he says that one can *try* to use ‘good’ or ‘bad’ predicatively, but such attempts are ‘illegitimate’ (p. 34). One interpretation of this is that whenever we use ‘good’ predicatively, we end up saying something meaningless, something that cannot be either true or false. Another possible interpretation would be that predicative uses of good make claims to the effect that some object possesses some non-relational and objective value property (objective goodness), and since there are no such properties in the world, we end up saying something meaningful but (systematically) false. In neither case need there be any disagreement between Geach and the value subjectivist on this point. However, subjectivists and objectivists would probably disagree with Geach’s claim that ‘good’ and other value words are always used attributively and that ‘we cannot sensibly speak of a good or bad event [or] a good or bad thing to happen’ (p. 41). We can, I believe, and often do, speak in precisely this manner. Whether we are thereby saying something that can be objectively true is a different matter.

⁵⁵ Cf. Korsgaard (2014). She draws a distinction between good-for in the functional (or ‘motherly’) sense and good-for in the final sense. She suggests that ‘the final sense and the motherly sense of “good for you” mention the same set of facts, but from two different perspectives. From one of these perspectives, we view Alfred as a functional system, that is, an entity whose parts and lesser systems all contribute to the achievement of some end or ends, in some cases simply to the maintenance or continuation of that functional system itself. From the other, we view the things that are good for Alfred from Alfred’s own point of view’ (p. 417).

their continued existence. Nitrogen would be objectively valuable for such plants.⁵⁶

I believe that subjectivists and objectivists alike can agree that it is objectively the case that nitrogen helps to sustain the life of many varieties of plants, and thus that nitrogen is ‘objectively good’ for such plants in a descriptive sense, whilst at the same time deny that the statement thereby makes any reference to an objective value, or necessarily implies that there is a normative reason to provide plants with nitrogen or to have a certain attitude—indeed to do anything in particular. Subjectivists and objectivists alike can agree that, given the aim of sustaining a plant’s continued existence, it will be a thoroughly objective matter whether providing it with sufficient nitrogen promotes this aim or not. It does not follow from this that whether providing plants with nitrogen is good or something that we have reason to do is an objective matter. Hence, allowing that nitrogen is objectively good for plants (in the sense described by Moore) commits us to nothing that would be disputed by subjectivists and objectivists. Such usages of ‘good for’ can therefore be put to one side.

I will assume that there a different—an evaluative or normative—usage of ‘good for’ in which to say that something is good for a person (or for a non-human animal, plant, artefact, etc.) is to express a genuine value judgment. It may not always be easy in practice to distinguish evaluative from non-evaluative uses of ‘good for’, but insofar as this is a problem it is a problem for subjectivists and objectivists alike. Perhaps this notion of good-for should be seen as a thick value notion, as having a fixed descriptive content in addition to its evaluative content. If so, there would be certain descriptive constraints on how goodness-for can be assigned. Or perhaps it should instead be seen as a thin value notion.⁵⁷

⁵⁶ Moore (2004) (p. 81). This is intended as a counterexample to Mackie’s claim that a special faculty would be needed to be aware of objective values. Whilst Moore is certainly right that ‘[n]o special faculty of moral perception or intuition is needed to understand [that nitrogen would be objectively good for many varieties of plants]’ (ibid), this fails as a counterexample if the usage of ‘good for’ is descriptive and not evaluative.

⁵⁷ Some philosophers think that the notion of good-for is synonymous with welfare or well-being. As Rønnow-Rasmussen (2011) has pointed out, this might be problematic if welfare is treated as a thick value notion and good-for as a thin value notion (pp. 106–107). How, after all, can a thick notion be synonymous with a thin one? If welfare is not a thick evaluative

Goodness-for is an example of what I will call a relational value, whilst straightforward goodness, or goodness *simpliciter*⁵⁸ is a non-relational value. Some philosophers hold that there are both of these kinds of values whilst others hold that there is only one of these kinds. Toni Rønnow-Rasmussen has proposed the following classification of positions on this matter.⁵⁹ Goodness-type dualism recognizes both relational and non-relational values and claims that neither of these types of value is reducible to the other. Goodness-type monists deny this and hold instead that one type of value can be fully understood in terms of the other.⁶⁰ Goodness-type monism comes in two main varieties. The first—Moorean monism—claims that all values are non-relational: all goodness is straightforward goodness or can be fully understood in terms of or reduced to straightforward goodness. The second—Hobbesian monism—claims that all values are relational: all goodness is goodness-for or can be reduced to goodness-for.⁶¹ The divide

notion but rather not an evaluative notion at all, it is even harder to see how goodness-for can be understood in terms of welfare. If both good for and welfare are treated as thin notions, synonymy can be maintained. But is it really plausible to treat welfare as a thin notion? This may well be doubted. Moreover, we might be sceptical about the suggested synonymy in the first place. It is not unreasonable to insist that a person's welfare is just one out of the many different things that can be good for the person. Admittedly, saying that a person's welfare might not be good for him or her does sound odd. But this oddity might be due merely to the fact that it militates against our first order intuitions about what is good for people, not due to it being conceptually confused.

⁵⁸ David Gauthier (1986) contrasts 'goodness for' with 'straightforward goodness' (p. 50). Rønnow-Rasmussen (2011, 2013) contrasts 'good-for' with 'good period', as does Kraut (2011). Kraut adds 'Latin and French phrases can be used to the same effect: we can ask whether we should value certain things because they are good *simpliciter*, *tout court*, *sans phrase*' (p. 5). At one point Kraut also speaks of 'mere goodness' (p. 8). It is also common to use the term *absolute* to refer to what I am calling non-relational values, but this usage will be avoided as it invites an illegitimate conflation of relational/non-relational with relative/absolute, two distinctions I will argue are important to keep apart.

⁵⁹ Rønnow-Rasmussen (2013)

⁶⁰ Note that goodness-type monism should not be confused with either axiological monism (the view that there is only one final value) or with value-bearer monism (the view that all bearers of final value belong to one and the same ontological category).

⁶¹ Naming this view after Thomas Hobbes was inspired by his claim that 'one cannot speak of something as being simply good; since whatsoever is good, is good for someone or other'.

between these forms of dualism and monism is independent of that between subjectivism and objectivism, and hence both dualism and the two distinct forms of monism are available to subjectivists and objectivists alike.

Just as subjectivists and objectivists do not offer answers to the question of what is good or bad, they do also not offer answers to what is good or bad *for* people (animals, plants, etc.). These are first order questions. But insofar as they recognize both types of value (that is, insofar as they are goodness-type dualists), subjectivists and objectivists will have to offer analyses of both goodness and goodness-for; they will have to be able to provide analyses of the nature of both relational and non-relational values. A subjectivist analysis of goodness-for will be sketched in Chapter 6.

IX. A Brief Roadmap

In this chapter I have characterized the philosophical project in which I take subjectivists and objectivists alike to be engaged. A guiding assumption has been that subjectivism and objectivism should be seen as rival analyses of the nature of value and that such analyses should not be confused with any first order (normative or axiological) theories or with any semantic analyses of evaluative language. As noted at the outset, in order to be rivals, subjectivists and objectivists must be in the business of analysing (at least partly) one and the same thing; subjectivist and objectivist analyses must have a common (or at least partly overlapping) analysandum. I do not think that I have fully defused the worry that subjectivists and objectivists might be talking about different things and thus are systematically talking past each other. But I have listed a few common points on which I think subjectivists and objectivists may be able to agree.

De Homine (XI.4.). As I will argue in Chapter 6, this claim is open to more than one interpretation, and in light of this it is not obvious that Hobbes himself was a Hobbesian monist.

Among these points we find the ideas that value depends/supervenes on the non-evaluative features of value bearers, that for something to be finally valuable is for it to be valuable for its own sake, that the intrinsic value of something depends/supervenes solely on the internal features of the object, and that whenever there is value there are reasons to respond to the value bearer in some way or other—perhaps by taking action towards it or by having some attitude towards it. I have insisted that both subjectivists and objectivists are entitled to talk about such things as value supervenience, final and intrinsic value, and reasons to respond to valuable objects, though of course—as will become clearer in the chapters that follow—they have very different conceptions of the nature of these things.

In the next chapter I identify what is distinctive about the subjectivist position in value theory. The subjectivist analysis of the nature of value is characterized as the view that value is grounded on attitudes. This I distinguish from certain other views with which it should not be confused. I will describe how subjectivists can analyse intrinsic value and respond to a couple of concerns about this way of understanding subjectivism.

In Chapter 3 I discuss how value subjectivists should deal with hypothetical attitudes and attitudes concerning hypothetical situations. I consider some objections against subjectivism that can be understood as asking us to imagine situations in which we had different attitudes than those that we in fact have. Insofar as our attitudes would be different from what they actually are, it seems that what is good and bad would be different also. This presents a problem, given that value is not supposed to behave in such a way; value is supposed to be invariant. I argue that subjectivist theories respect the relevant invariance. This is because it is always the attitudes that we actually have that ground values, even in hypothetical situations. Another desideratum is consistency, which I argue that subjectivists can also respect. There may, however, be more than one way to go about it.

Chapter 4 is devoted to two interrelated topics: motivation and idealization. I argue that the attitudes on which values are grounded are specifically conative attitudes—attitudes that involve a motivational component. A special form of motivational internalism (which I call perspective internalism) is linked with the value subjectivist view. The second main objective of the chapter is to consider ways in which a subjectivist might want to further restrict the class of value-grounding attitudes, or otherwise idealise these attitudes. I discuss which (if any) forms of idealization that can be accepted by a value subjectivist.

In Chapter 5 I distinguish two different versions of value subjectivism on the basis of whose attitudes that are taken to ground value: is it *my* attitudes (those of the individual subject), or is it instead *our* attitudes (those of all subjects somehow taken together)? I characterize these two different versions of value subjectivism, and, though I ultimately leave the debate between them in an unresolved state, I argue that there is a presumption in favour of the former (individualist relativist) version of subjectivism over the latter (universalist absolutist) version.

The sixth chapter is devoted to the topic of relational values, such as goodness-for. Insofar as a value subjectivist is not a Moorean monist but either a Hobbesian monist or a goodness-type dualist, the question arises of how a subjectivist might analyse relational value. I sketch one such analysis.

In Chapter 7 I discuss how we might understand 'projection' of value and then proceed to discuss the implications of projectivism for the ontological status of value. The most plausible stance for a subjectivist to take is anti-realist. In the eighth chapter, I discuss which theories of evaluative language might be accepted in conjunction with a subjectivist theory of the nature of value. Finally, there is a brief afterword with some concluding thoughts.

2. VALUE GROUNDED ON ATTITUDES: SUBJECTIVISM IN VALUE THEORY

[T]he way in which moral entities are produced can scarcely be better expressed than by using the word *imposition*. For they do not arise out of the intrinsic nature of the physical properties of things, but they are superadded, at the will of intelligent entities, to things already existent and physically complete.

Samuel Pufendorf, *The Law of Nature and of Nations*

I. Introduction

It is commonly thought that what differentiates subjectivists from objectivists with respect to value is that, whilst subjectivists take value to be related to pro and con attitudes, such as desires or preferences⁶², objectivists deny any such relation. But that values and attitudes are in some way related, even necessarily so, is not something that objectivists need to deny. Consider, for example, the objectivist version(s) of the fitting-attitude analysis of value mentioned in Chapter 1. On such views, for something to

⁶² Recall that I am using the terms *pro* and *con attitude*, and the corresponding terms *favouring* and *disfavouring*, as placeholders for a wide range of different attitudes, including but not limited to desire and preference.

be valuable is for it to be the object of an attitude that it is objectively *fitting* to have, or that there is an objective *normative reason* to have. Though such views take value to be necessarily related to attitudes, they are essentially objectivistic in that the standard of fittingness, or the reason to have a certain attitude, is not grounded on attitudes.

An even less accurate way of drawing a line between objectivism and subjectivism would be to say that, on subjectivist views, value *depends* on attitudes, while on objectivist views it does not so depend. This is unsatisfactory as there is an important sense of *depend* in which a subjectivist may deny that value necessarily depends on attitudes and some objectivists may affirm a necessary dependence of values on attitudes in the same sense. As will be seen, what it is for value to *depend* on attitudes is open to radically different interpretations, not all of which are best seen as subjectivist. I prefer to use the term *depend* to refer to the relation of supervenience, the dependency relation that holds between the value of an object and the non-evaluative features of the object in virtue of which it has its value (plus enablers, disablers, etc., insofar as these are recognized). When I henceforth talk about value depending on something, or about an object having value in virtue of something, or about some set of non-evaluative properties making an object good or bad, I am talking about the supervenience relation. I will argue that value subjectivism is not the view that value depends (in the sense of supervenes) on attitudes, that all valuable things have their values in virtue of being objects of attitudes, or that what makes an object good or bad is always the fact that it is favoured or disfavoured by some subject. That would be a first order axiological view that could be accepted or rejected by subjectivists and objectivists alike. We need to make a distinction between dependence (supervenience) and the fundamentally distinct relation of grounding or constitution. Subjectivism, I argue, is the view that value is grounded on or constituted by attitudes.

II. Backgrounding Attitudes

When asked, with respect to something you value, ‘what makes this thing good, or better than some other thing?’ your answer will seldom take the form of some fact about yourself, such as that you like the object in question

or that you prefer it over the other. Instead, your answer will usually be in terms of something about that object. Specifically, your answer would commonly take the form of a list of the features of the object *for which* you value it. This simple observation may seem to constitute a problem for any view that wishes to understand value in terms of attitudes. If what we generally think makes a thing good is something about that thing (perhaps the way that thing is in itself) and not the fact that it is desired or preferred by anyone, then values might seem to be sufficiently independent of attitudes to be properly considered 'objective'. However, that value is ascribed to objects and that their value somehow depends (supervenes) on the non-evaluative features of the objects in question is not something to which value objectivists can claim a monopoly. Indeed, as I insisted in the previous chapter, this is common ground among subjectivists and objectivists; it forms part of the very analysandum of which they offer rival analyses.

The first thing I wish to draw attention to here is that attitudes can figure in the background of valuing without figuring in the foreground. This claim has been defended by Philip Pettit and Michael Smith as a more general claim about motivation. They argue that desire can have 'a motivational presence in my decision making, figuring in the background, as it were, without appearing in the content of my deliberation, in the foreground'.⁶³ Stephen Darwall makes use of the same distinction when he says 'Desires are *for* their objects, and the locus of the desiring person's concern is most typically the object itself and the facts connected with it rather than the fact that she has the desire'.⁶⁴ In an essay about Thomas Hobbes' theory of value, Darwall says similarly that (according to Hobbes)

[a]ll deliberation begins in an agent's desires, but this does not mean that they begin in a premise *about* her desires. The deliberating agent reasons from a premise she accepts *in* having a desire, not from the premise *that* she has a desire.⁶⁵

⁶³ Pettit & Smith (1990) (p. 565)

⁶⁴ Darwall (1983) (p. 37)

⁶⁵ Darwall (2000) (p. 333). The emphases in this and the previous quote are mine.

Whilst some desire (or other pro or con attitude) is always present when we act and can be appealed to in an explanation of the action⁶⁶, we are seldom motivated to act by our awareness of the fact that we desire something. What motivates us is rather the object or something about the object of our desire. Simon Blackburn puts it well when he says ‘From the inside, the objects of our passions are their *immediate* objects: it is the death, the loved one, the sunset, that matters to us. It is not our own state of satisfaction or pleasure.’⁶⁷ As Blackburn also points out, it is a mistake to suppose ‘that when we deliberate in the light of various features of a situation we are *at the same time* or “really” deliberating—or that our reasoning can be “modelled” by representing us as deliberating—about our own conative functioning.’⁶⁸ As Pettit and Smith recognize, the deliberating agent may sometimes ‘take cognizance of the fact that he has this or that passion or yen or hankering. But such self-concern seems to be the exception, not the rule’.⁶⁹

Consider also the following example formulated by Wlodek Rabinowicz. Suppose that I have a desire to understand the theory of relativity. I want to achieve this understanding for its own sake, out of pure intellectual curiosity. If I were to achieve this understanding my desire would thereby be satisfied, but what I want to achieve for its own sake is the understanding (that is, the object of my desire) and not the satisfaction of my desire (whatever its content). The satisfaction of my desire need not be something that I want to realize, at least not for its own sake.

⁶⁶ As Darwall (1983) writes: ‘From the fact that any intentional action must be explainable by the agent’s desires it does not follow that the agent’s reasons for acting must itself make reference to his desires’ (p. 37). Note that the value subjectivist is not necessarily committed to any particular view in the philosophy of action. However, at least one philosopher has argued that it is a virtue of the subjectivist (what he calls ‘projectivist’) view that it is readily compatible with the ‘standard model of explanation of why someone does something’. This model attributes both a belief and a desire to the agent: ‘The belief that a bottle contains poison does not by itself explain why someone avoids it; the belief coupled with the normal desire to avoid harm does’. The argument is spelled out in Blackburn (1984) (pp. 187–189). Motivation will be further discussed in Chapter 4.

⁶⁷ Blackburn (1988) (p. 370)

⁶⁸ *Ibid* (p. 371)

⁶⁹ Pettit & Smith (1990) (p. 575)

I *might* want it, and I probably do, along with lots of other things that I desire, but it is not necessary that I have this ‘second order’ desire to have my first order desire satisfied. In particular, it may well be the case that my desire to understand the relativity theory is ‘categorical’ in the sense that the object of my desire is to understand the theory in question and not ‘to understand it *provided* that I desire to understand it.’ In other words, my desire need not be conditioned on its own existence. But then I *might* be indifferent to the fact that, by bringing about the object of my desire, I also bring about that my desire is being satisfied. The satisfaction of my preference need not be one of the objectives I am motivated by.⁷⁰

My desire for ice cream, by contrast, is a desire that is indeed ‘conditioned on its own existence’.⁷¹ I desire ice cream *only provided that I desire it*. When I imagine a situation in which I don’t desire ice cream, I fail to have any positive attitude towards me having ice cream in the imagined situation. In this case, my *desire* for ice cream (and not just the ice cream itself, or the fact that I’m having it) does figure in the foreground of my deliberation. The point I presently want to make is that not all (indeed few) of our desires seem to be like my desire for ice cream in this sense, and more like the ‘categorical’ desire to understand the theory of relativity.⁷²

⁷⁰ Rabinowicz & Österberg (1996) (p. 7)

⁷¹ Note that this is not the same as Derek Parfit’s expression ‘conditioned on its own *persistence*’. See Parfit (1984) (p. 157).

⁷² Mark Schroeder (2007b) appeals to the distinction between desires figuring in the background and desires figuring in the foreground in his theory of normative reasons. On Schroeder’s view, desires are ‘background conditions’ on reasons; desires help explain why facts (or true propositions) are reasons, but desires need not themselves be *part of* the reasons (pp. 29–30). To borrow two of Schroeder’s examples, when Ronnie is asked what his reason was for going to the party or what his reason was for helping Katie, the most natural answers would be ‘that there was dancing at the party’ or ‘that Katie needed help’, none of which make explicit reference to desires. This is exactly what we should expect if desires do not always figure in the foreground. But it does nothing to undermine the view that Ronnie’s desires figure in the background—that they are background conditions on his reasons. When asked about his reason to go to the party, it might be equally natural for Ronnie to answer ‘that I wanted to dance’. In that case, Ronnie’s desire for dancing is like mine for ice cream—a desire conditioned on its own existence. In the case of helping Katie, on the other hand, a reply that made explicit reference to Ronnie’s desire to help would strike us as odd. Again, this should not surprise us given that desires do not always figure in the foreground. It does nothing to shed doubt on the view that Ronnie’s desires figure in the background. It seems

When you deliberate about how to act, you generally decide to do what you do because you think the course of action in question has some feature(s) other than the feature of *being desired by you*. There is always something about the option in virtue of which you choose it, and this something is rarely the very fact that you desire it. The premises of your practical reasoning mention the various characteristics of possible actions towards which you have some positive or negative attitude. But the fact that you have a certain attitude towards a course of action is generally not part of why you choose that course of action, and it is generally not what makes that course of action a good or a bad one in your eyes.⁷³ As Blackburn puts it, ‘What makes cruelty abhorrent is not that it offends us, but all those hideous things that make it do so’.⁷⁴ None of this should be surprising given that desires do not always appear in the foreground, in the content of the premises of practical reasoning. However, as Pettit and Smith insist, that a desire does not appear in the foreground of deliberation is consistent with the claim that it does appear in the background.

These authors also note that a desire that appears only in the background need not be ‘phenomenologically salient’.⁷⁵ That it does not need to be so might be part of the explanation for the belief that desires do

then that neither in our deliberation, nor in our justifications, need desires figure in the foreground.

⁷³ Pettit & Smith (1990) (p. 579)

⁷⁴ Blackburn (1988) (p. 367)

⁷⁵ Pettit & Smith (1990) (pp. 579–580). The distinction between a desire that appears only in the background and a desire that appears also in the foreground should not be identified with David Hume’s distinction between ‘calm’ and ‘violent’ passions. The former distinction concerns whether or not a desire figures in the content of deliberation. The latter distinction concerns the nature of the felt sensation that a desire produces in the agent. As Hume wrote in the *Treatise*, ‘there are certain calm desires and tendencies, which, tho’ they be real passions, produce little emotion in the mind, and are more known by their effects than by the immediate feeling or sensation’. ‘Beside these calm passions, which often determine the will, there are certain violent emotions . . . which have likewise a great influence on that faculty . . . [but that] rise to a great height, and produce a sensible emotion’ (ii.iii.iii.). Hume argues that calm passions are easily confused with beliefs, and as such can give rise to the mistaken view that it is through the faculty of reason (not passion) that we distinguish between good and bad (vice and virtue). The two distinctions seem to be orthogonal to each other.

not have to enter the picture at all.⁷⁶ But, as the same authors note, this belief gains no support from the phenomenology of deliberation as this phenomenology shows only that desires do not always figure in the foreground of deliberation—and does nothing to undermine the view that desires always figure in the background.

What Pettit and Smith call ‘the strict background view of desire’—that desires always figure in the background but not always in the foreground—is itself neutral between subjectivism and objectivism.⁷⁷ But recognizing that desires can figure in the background without figuring in the foreground dissolves a common misunderstanding about subjectivism by helping us to see that the subjectivist is by no means committed to the implausible claim that whenever we talk and think about what is good and bad we must necessarily talk and think about desires or other attitudes. The value subjectivist is not committed to saying that attitudes must always figure in the foreground of valuing and neither is the objectivist committed to denying it. Indeed, it is consistent with subjectivism that they *never* do. All that subjectivists need to say is that attitudes always figure in the background of valuing. This is why the simple observation with which I started—that what is held to make something good is normally something about that thing and not something about the evaluator—does not pose any obvious problem for the subjectivist.

⁷⁶ A rather extreme view of this type is held by Jonathan Dancy (2000).

⁷⁷ Pettit and Smith (1990) point out that the strict background view of desire (and the argument that they give for this view) is neutral between cognitivism and non-cognitivism. However, it might not be compatible with every form of cognitivism. Pettit and Smith think that it ‘rules out the cognitivist position according to which desire is not needed for someone to act on a reason. But it does not rule out cognitivism as such’ (p. 574). See also (*ibid.*, pp. 579–580).

III. Two Interpretations of Preferentialism in Value Theory

The idea that value is determined by attitudes is open to two radically different interpretations.⁷⁸ In a joint paper focusing on preferences, Wlodek Rabinowicz defends the object interpretation and Jan Österberg defends the satisfaction interpretation of preferentialism in value theory. On the satisfaction interpretation, final value is ascribed to the circumstance *that* our preferences be satisfied. On the object interpretation, by contrast, final value is ascribed to those things or states of affairs that are the *objects* of our final preferences—that is, to that which we prefer for its own sake.⁷⁹ I will argue that it is the latter interpretation that should be identified with subjectivism.

Rabinowicz and Österberg give us the following illustration of the difference between the two interpretations:

When we wish, say, that rainforests survive (not, or at least not only, [instrumentally], but for their own sake), is it then [finally] valuable that we get what we wish [as the satisfaction interpretation would have it], or is it rather the survival of the rainforest that, in the light of our wish, presents

⁷⁸ I have chosen the expression *determined by* as neutral between the two interpretations. This means that adherents of both interpretations hold that value is determined by attitudes. This should be kept distinct from the claim that value is *grounded* on attitudes, which is made by the object interpretation alone.

⁷⁹ As noted in Chapter 1, if subjectivism is formulated in terms of preferences alone, it is perhaps natural to think of the objects of value as being exclusively states of affairs, and thus to think of the subjectivist as adhering to a version of value-bearer monism. But insofar as subjectivists are prepared to allow a wider set of pro and con attitudes, including, for example, admiration, adoration, cherishing, honouring, respecting, loving, etc., they should probably also be willing to allow that value may accrue to objects belonging to several different ontological categories, and thus endorse a pluralistic view of value bearers.

itself as something [finally] valuable [as the object interpretation would have it]:⁸⁰

We get what we wish if and only if the rainforest survives. But there are fundamental differences between the two interpretations. On the satisfaction view, value accrues to the compound state consisting of a preference and the object of that preference being realized. That is, final value always accrues in part to preferences.⁸¹ On the object view, by contrast, final value is ascribed exclusively to the objects of preferences (in virtue of the non-evaluative properties of the objects in question *for which* they are being preferred) and not even in part to the preference itself (except in those cases where the preference itself forms part of the object of a preference—where the

⁸⁰ Rabinowicz & Österberg (1996) (p. 2). I have substituted the original ‘intrinsic’ here with ‘final’ value, and I will continue to make such substitutions whenever appropriate throughout this work. This rests on the conviction (explained in Chapter 1) that the final value of something need not be intrinsic. Rabinowicz and Österberg acknowledge this possibility in a footnote (*ibid.*, p. 3, n. 1) and, as previously mentioned, Rabinowicz has later defended the claim that final values need not be intrinsic in Rabinowicz & Rønnow-Rasmussen (2000).

⁸¹ The satisfaction view should be kept distinct from axiological hedonism, the view on which what is finally valuable is not the satisfaction of preferences but the feeling of satisfaction, or pleasure that one normally gets from having one’s preferences satisfied. However, though a distinct position in its own right, the satisfaction view is open to some of the same objections that have been aimed at axiological hedonism. Rabinowicz (in Rabinowicz & Österberg (1996)) applies Joseph Butler’s criticism of hedonism to the satisfaction interpretation: ‘[P]reference satisfaction is only possible as a secondary aim; if our only preference were the desire to have our preferences satisfied, we would then have no preferences to satisfy (apart from this second order preference, which would then be vacuously satisfied)’ (p. 8). In Butler’s terminology of the Preface to *Fifteen Sermons* ‘self-love’ needs ‘particular passions and appetites’ as input. In the absence of such input, it would have ‘absolutely nothing at all to employ itself about; no end or object for it to pursue’ (Section 31). According to Rabinowicz (*ibid.*), an axiological view on which nothing but preference satisfaction has final value similarly ‘deprives our [final] preferences of their authority by denying their objects [final] value, while at the same time it feeds on these denigrated preferences.’ This position ‘sponges on the attitudes that it in the same breath disavows’ (p. 9). The satisfaction interpretation does thus not seem to take our preferences seriously, though it seems to presuppose that *we* (as preferring subjects with ‘particular passions and appetites’) continue to take them seriously. My aim is not to argue against the satisfaction view as such, but only to distinguish it from subjectivism.

satisfaction of the desire is itself desired, exemplified above by my desire for ice cream).

It should be noted that the object interpretation goes very naturally together with the strict background view of desire discussed in the previous section. Recall the example of wanting to understand the theory of relativity for its own sake. What motivates me in that case is achieving the understanding (that is, the object of my desire) and not (at least not exclusively or even primarily) the satisfaction of my desire to understand. This fits very neatly with the object interpretation on which final value is ascribed precisely to that which is finally desired. On this view, as Rabinowicz emphasizes, 'the pertinent [final] value is precisely the objective that I am motivated by: the object of my desire.'⁸²

The two interpretations carry interesting implications. On the satisfaction interpretation, the only thing that is valuable for its own sake is preference satisfaction and anything else carries only non-final value. Ascription of final value to things other than preference satisfaction (such as the understanding of a theory or the survival of a rainforest) is on this view a systematic mistake. On the object view, by contrast, anything that can be the object of a final attitude can be an object of final value. In the words of David Gauthier, whom I take to be a defender of the object interpretation⁸³, 'there is no restriction on the nature of those states of affairs that may be objects of preference, and so may be valued'.⁸⁴ And, as Rabinowicz writes, since our final preferences 'seem to have greatly heterogeneous objects', the object interpretation 'naturally gives rise to a plethora of values'.⁸⁵ Whilst the satisfaction interpretation is a version of axiological monism, the object

⁸² Rabinowicz & Österberg (1996) (p. 7)

⁸³ Whilst he does not apply the labels *object interpretation* and *satisfaction interpretation* (which were coined ten years later by Rabinowicz and Österberg), it is clear that Gauthier (1986) adheres to the object interpretation: '[V]alues are ascribed to states of affairs', he says, and 'the ascription is attitudinal, not observational, subjective, not objective' (p. 25). Like Rabinowicz & Österberg (1996), Gauthier too formulates the thesis in terms of preferences for states of affairs, but as previously noted, the relevant attitudes need not be restricted to preferences and value bearers need not be limited to states of affairs.

⁸⁴ Gauthier (1986) (p. 47)

⁸⁵ Rabinowicz & Österberg (1996) (p. 11)

interpretation leads instead to a pluralistic axiology according to which 'there are several irreducibly distinct [final] values'.⁸⁶

The most interesting contrast for my purposes is this: on the satisfaction interpretation, the sole thing that carries final value (preference satisfaction) is taken to have its value irrespective of whether it is being the object of an attitude; preference satisfaction is held to be finally valuable whether or not it is finally preferred. In this, I argue, defenders of the satisfaction interpretation hold preference satisfaction to be *objectively* valuable. Hence, even though the satisfaction interpretation makes essential reference to preferences, it is not a subjectivist view. However, there is a way that defenders of the satisfaction interpretation might be seen as avoiding commitment to objectivism. They could say that, since the satisfaction view is a first order (axiological) view—a view that tells us what is valuable—it can remain silent on the second order issue of the objective or subjective nature of value. Since this is so, the two views are strictly speaking compatible with each other.⁸⁷ Indeed, under the assumption that our final preferences invariably take preference satisfaction as their object, the object interpretation would seem to imply the satisfaction interpretation! But as I have already insisted, this assumption is empirically implausible. Indeed, not only is it not always so, but it seems seldom to be the case that what we want for its own sake is only that our preferences be satisfied.⁸⁸

We might ask how the object interpretation, in virtue of having been linked to axiological pluralism, can be compatible with the monistic

⁸⁶ Ibid

⁸⁷ Björn Petersson has pointed out the compatibility of the two interpretations to me. Of course, as *interpretations* considered, the two views are not compatible; interpreting value preferentialism in one of these ways precludes us from at the same time interpreting it in the other way. But if we are willing to see the satisfaction interpretation as a first order view and the object interpretation as a second order view, then the two views are strictly speaking compatible.

⁸⁸ As Gauthier (1986) says, there 'may be some persons who reveal and express preferences only in relation to their own experience, but there are surely few whose preferences are so confined' (p. 47). We may say the same about preference satisfaction: there are surely few, if any, who desire only the satisfaction of their preferences for its own sake. Furthermore, as was pointed out in footnote 81 above, this may also be problematic in its own right in light of Joseph Butler's criticism of hedonism (applied to the satisfaction interpretation of preferentialism).

satisfaction interpretation. The answer is, I believe, that whilst the object interpretation gives rise to, or leads to, a pluralistic axiology when applied to a world like ours, in which there exist final preferences with ‘greatly heterogeneous objects’, it does not in and of itself entail any first order axiological view. I think we should say that whilst the object interpretation is a subjectivist second order view with pluralistic first order implications, the satisfaction view is a monistic first order view with objectivist second order implications.

As stated in Chapter 1, I treat subjectivism and objectivism as competing second order theories: as rival analyses of value. My purpose is thus not to argue against the satisfaction view as such, or to argue for or against any first order view, but only to distinguish such views from subjectivism. This is important, since many writers who have taken themselves to be criticizing subjectivism have in fact been criticizing something like the satisfaction view. Or worse, failing to recognize the distinction, oscillating between criticizing the satisfaction view (or similar views) and views that can properly be considered versions of subjectivism in my sense, all lumped together under the same heading. Adam D. Moore seems to be guilty of such a conflation when he writes,

People value a wide range of objects, activities, goals, careers, and pursuits. When asked what is valuable we include things like a nice day on the golf course, hanging with friends, spicy Indian food, a fast car, lots of money, and good music. This list could be continued. But what thread runs through each of these items that make it count as good? One common answer is that they are good because they are desired. On this view what is [finally] good, or good in and of itself, is what satisfies a [final] desire. . . We desire such things and when they come to pass value has been brought into the world. Such views about value are commonly called subjective preference satisfaction theories or desire satisfaction accounts.⁸⁹

Moore’s project is to show, by way of a series of counterexamples constructed to engage our first order intuitions, ‘the implausibility of maintaining the claim that the sole standard of value, in fact, that which

⁸⁹ Moore (2004) (pp. 75–76)

creates value, is the satisfaction of desires and preferences'.⁹⁰ This characterization seemingly runs elements of the object and satisfaction interpretations together and thus implicitly assumes that arguments against one interpretation will also be arguments against the other. But, in fact, since the two interpretations are fundamentally different, arguments that might be appropriate and effective against one interpretation might leave the other interpretation completely unscathed (and vice versa).

The idea that things are good because they are desired is open to two radically different interpretations. We have already seen that when asked what makes something—a nice day on the golf course, hanging out with friends, or spicy Indian food, (to appropriate some of Moore's examples)—good, we should not in general expect the answer to have anything to do with desires or preferences. Whilst it is indeed the case on the subjectivist view that value is taken to be created or brought into the world by attitudes—a claim perspicuously not made by the satisfaction interpretation—the subjectivist is by no means committed to a first order view that takes desires and preferences or their satisfaction to be the sole standard of value. On the contrary, as we have seen, the object interpretation seems to give rise to a pluralistic axiology on which many different things—none of which need necessarily involve desires or preferences—may be objects of final value.

IV. Ground and Object

I have maintained that subjectivism is best seen as a second order view and that it should not be confused with the first order view that only preference satisfaction is finally good, or with any other first order view. Other authors have drawn this same conclusion. David Gauthier insists that we must not suppose 'that subjectivism implies that only what it is itself subjective—only a state of sentient experience—could possess [final] value'.⁹¹ Toni Rønnow-

⁹⁰ Ibid (p. 76)

⁹¹ Gauthier (1986) (p. 47)

Rasmussen similarly urges us to move beyond ‘one-dimensional approaches that regard . . . a subjectivist as someone who locates value to subjective states’ and to recognize that ‘subjectivists and objectivists may well share each other’s evaluations’.⁹² Gauthier too maintains that subjectivism and objectivism are in fact ‘equally neutral with respect to the nature of the objects of [final] value’, and that objectivism is ‘quite compatible with, although certainly it does not require, the view that the only bearers of [final] value are states of sentient experience’ (as on axiological hedonism, for example). In more general terms, Gauthier says that ‘the difference between subjective and objective conceptions of value is not a difference between the objects or states of affairs to which value may be ascribed’. The debate between subjectivists and objectivists ‘concerns the ground of value, not its object’.⁹³

The distinction between ground and object will loom large in the present work. As I have previously hinted, even though attitudes can play an essential role in an objectivist theory of value, subjectivists make a claim that objectivists deny: namely, that value is created or brought into the world by attitudes. Or, as Gauthier puts it, though objectivists need not deny that values are ‘affectively related’ (as they clearly are on the satisfaction interpretation, for instance), they do deny that values are ‘products of such relationships’. On subjectivist views, by contrast, ‘values are products of our affections’.⁹⁴ This is a claim about the ground of value. As such, it should not be understood as in any way conflicting with the very different claim that the value of an object somehow depends (supervenes) on its non-evaluative features. This claim, as I argued in Chapter 1, is common ground among subjectivists and objectivists. In order to understand what is distinctive about subjectivism, it is of the essence that we do not conflate claims about the ground of value with claims about the object(s) of value.

As I noted in Chapter 1, subjectivists and objectivists agree that value is something that accrues to objects and that it supervenes on the non-evaluative features of these objects. Subjectivists hold, in addition, that whatever value an object has, this value is conferred upon it by some attitude

⁹² Rønnow-Rasmussen (2003) (p. 247)

⁹³ Gauthier (1986) (p. 47)

⁹⁴ *Ibid.* My emphases.

that some subject has toward the object in question, and the value is taken to supervene on precisely those non-evaluative properties of the object *for which* it is desired or preferred.⁹⁵ As Rønnow-Rasmussen writes, that an object is valuable is, on subjectivist views, grounded ‘in something *external* to the object; it is the pro- et contra attitudes of the subject that bestow value on the object towards which they are directed.’⁹⁶ Value is thus seen as coming into being by being conferred upon otherwise valueless objects by the attitudes that we have towards the objects in question.

The Notion of a Constitutive Ground of Value

It should be noted that the distinction between ground and object underlies the object interpretation discussed in the previous section. In the same paper in which the object and satisfaction interpretations are played out against each other, Rabinowicz introduces the technical notion of *constitutive ground*. He writes ‘[we] need to distinguish between the features on which [final] value *supervenes* and the features that are constitutive grounds of [final] value’.⁹⁷

According to the object interpretation, [final] values of states of affairs are in a sense created by something external—by our ([final]) preferences, but they do not supervene upon the preferences. Instead, they supervene on those . . . properties of the states of affairs for which the states in question are being [finally] preferred.⁹⁸

Value is thus taken to be created or constituted by the attitudes of subjects.⁹⁹ I take this claim—that value is constituted by subjects—to be the defining

⁹⁵ Insofar as we allow persons as possible value bearers, the object of value might be the subject herself.

⁹⁶ Rønnow-Rasmussen (2003) (p. 251)

⁹⁷ Rabinowicz & Österberg (1996) (pp. 20–21)

⁹⁸ *Ibid* (p. 21)

⁹⁹ This usage of *constitution* should not be confused with the different usage of the term in which being constituted by something means that it *consists thereof* that it is *part* of it. To better bring out the present usage, focus could be put on the verb form ‘to constitute’. It has

claim of value subjectivism. Rønnow-Rasmussen reaches the same conclusion, arguing that the difference between subjectivism and objectivism is that 'subjectivists make a claim that objectivists deny or at least interpret differently, viz. that *values are constituted by subjects*.'¹⁰⁰

V. How Subjectivists Can Analyse Intrinsic

(Final) Value

Saying that the value of an object is constituted or created by something external to the object itself (by an attitude that some subject is directing towards the object) might seem to lead us towards the conclusion that subjectivists cannot talk about intrinsic value.¹⁰¹ But this conclusion does not follow. It becomes clear that subjectivists can recognize intrinsic values

been suggested to me that, as an alternative to 'constitutive grounds', 'projective grounds' could be used. This would perhaps more clearly capture the nature of the subjectivist idea, but I will go along with the by now established constitutive ground. Rabinowicz tells me in personal communication that he was inspired by Husserl's transcendental phenomenology when he decided to use the term *constitution*.

¹⁰⁰ Rønnow-Rasmussen (2003) (p. 251). The fuller passage reads as follows: '[I]f claims about supervenience bases are evaluative claims, it cannot be in this neighbourhood that we should look for the difference between subjectivism and objectivism—at least not if we want to retain the value neutral metaethical character of these views. The difference lies elsewhere; subjectivists make a claim that objectivists deny or at least interpret differently, viz. that *values are constituted by subjects*.'

¹⁰¹ As noted in Chapter 1, some philosophers have held that only objectivists can talk about intrinsic value. G. E. Moore (1922), for example, took the alleged failure of the subjectivist to recognize intrinsic values to be the 'strongest ground of [the] objection to any subjective view' (p. 255). I will here argue that this belief rests on a failure to appreciate the distinction between supervenience base and constitutive ground. In this I follow Rabinowicz & Rønnow-Rasmussen (2000).

once we distinguish between supervenience base and constitutive ground. It is only the former—the features of the valuable object in virtue of which it has its value—that need to be internal to the object in order for the value to be intrinsic.

A value is intrinsic if it supervenes exclusively on the internal features of the valuable object. This much, as I insisted in Chapter 1, may be taken to be common ground among subjectivists and objectivists. On subjectivism, an object is intrinsically valuable when the features of the object for which it is being desired or preferred or otherwise favoured are all internal to the object. That is, an object is intrinsically good insofar as it is favoured in virtue of its internal properties.

Here, then, is the subjectivist analysis we have arrived at. Final value is grounded on and constituted by final attitudes, attitudes of favouring or disfavouring something for its own sake. Positive values such as goodness are grounded on pro attitudes, and negative values such as badness on con attitudes. Final intrinsic values are grounded on and constituted by attitudes of favouring or disfavouring something for its own sake in virtue of its internal features. As I argued in Chapter 1, we need not suppose that all final values are intrinsic. According to subjectivism, an object is finally extrinsically valuable insofar as it is favoured or disfavoured for its own sake in virtue of its relational features. Recall the example I used of the guitar that once belonged to Steve Howe. The guitar may be desired for its own sake, but in virtue not of its internal features alone but in virtue of its relation to Howe of having belonged to him in the past. If this is so, the subjectivist concludes that the guitar is of positive final extrinsic value relative to the desirer in question. Someone else may favour the same guitar strictly in virtue of its internal features, in which case the guitar is intrinsically valuable relative to that person. Yet another person may favour it as a means for something else (perhaps solely in order to make money from selling it), in which case the guitar is of merely instrumental value relative to that third individual. This relativistic aspect on value subjectivism will be further discussed later on.

VI. Objections

On the view we have arrived at, value is taken to supervene exclusively on the objects of pro and con attitudes (specifically on the features of objects in virtue of which they are favoured or disfavoured by subjects) and not on the attitudes themselves, except in those cases where an attitude is itself part of the object of an attitude. Subjectivists thus insist that though value is created or constituted by attitudes, that an object is favoured or disfavoured should generally not itself be counted among the value-making properties of that object. Being favoured by some subject is generally not part of the supervenience base of value.

One might object that there is no need to introduce the novel distinction between supervenience base and constitutive ground. Couldn't the subjectivist just accept that the attitude that is thought to constitute a value is itself part of the supervenience base of the value? One problem with this proposal is that it would force on us a restrictive view according to which only certain objects could carry value: namely, only states of affairs that involve some subject having some attitude towards something.¹⁰² Such states are not ordinarily what we value, and it seems overly restrictive to insist that we confine ourselves to ascribing value to only such states.

Another, more fundamental, problem with the objection is that if value would supervene on (and only on) such states, these states would seem to have their values regardless of whether they are objects of attitudes. This would, I believe, be asking the subjectivist to accept objective values! Unless it can somehow be shown that the distinction between supervenience base and constitutive ground is somehow incoherent, not allowing subjectivists to formulate their view by drawing this distinction seems to me to beg the question against the subjectivist. That the notion of a constitutive ground of value remains obscure might be a fair criticism, but insisting that we do away with it and explain everything in terms of supervenience by baking the constitutive attitude into the supervenience base would be insisting that subjectivists adopt an objectivist analysis of the nature of value.

¹⁰² This objection is discussed in Rønnow-Rasmussen (2011) (pp. 14–15). Rønnow-Rasmussen responds to the objection in this first way. I add a more fundamental counter to the objection.

Johan Brännmark poses a different challenge to what he calls ‘Quasi-Moorean projectivism’, which is the alternative label under which he discusses subjectivist views understood in the way I have been describing. Brännmark calls it quasi-*Moorean* because ‘[i]n spite of the fact that it repudiates the strong objectivity ascribed to values by Moore, it still seeks to employ essential parts of the Moorean terminology’ such as the notion of intrinsic value. Brännmark also indicates that subjectivism might actually be able to provide us with ‘a new way of looking at the archetypal Moorean situation that is featured in his method of absolute isolation, one that does much to demystify the scene. Rather than being a peculiar epistemic device, it becomes a schema of the ontic relation between subject and value’.¹⁰³ The challenge is this: Brännmark asks in what sense we are justified in speaking of value as in any way grounded on such views. This is an entirely reasonable worry that I will try to address in this work. But what would constitute a satisfactory answer? Brännmark specifies what he demands by saying that ‘the question is whether there really is a metaphysical substance “value” which has some kind of “source” [the subject] from which it can be “conferred” [upon objects]’, or instead whether ‘it is just a metaphysical-sounding way of expressing something that is not metaphysical, but rather a psychological phenomenon’.¹⁰⁴ Why should understand the distinction as one between supervenience base and constitutive ground, rather than, as Brännmark suggestively proposes, between ‘appreciative base’ and ‘imaginary ground’, ‘a distinction between the features which the subject sees as good-making in the object, and the (psychological) features that make the subject see something nice in the object?’¹⁰⁵

Whilst it is perfectly reasonable to require that subjectivists tells us more about what ‘value constitution’ is supposed to be, demanding that they do this in a way that treats value as a ‘metaphysical substance’ seems to beg important questions. For all that I have said so far, a subjectivist might hold that value is a metaphysical substance that is somehow brought into existence when subjects have attitudes. However, as I will argue more fully later on, the most plausible versions of subjectivism are anti-realist, and

¹⁰³ Brännmark (2002) (p. 66)

¹⁰⁴ Ibid (p. 67)

¹⁰⁵ Ibid (p. 68)

philosophers who have defended value subjectivism are predominantly anti-realists. Gauthier is a case in point. He argues that value is ‘not an inherent characteristic of things or states of affairs, not something existing as part of the ontological furniture of the universe’.¹⁰⁶ Rabinowicz points out that ‘the object interpretation presupposes a subjectivist (“projectivist”) theory of value’ on which value is ‘not part of the mind-independent world but something that we project upon the world’.¹⁰⁷

But perhaps stressing anti-realism does more to exacerbate the problem than to resolve it. Brännmark points to an ‘obvious tension’ in the subjectivist/projectivist theory: ‘On the one hand, there is a desire to present a scientifically respectable ontology; on the other hand, the desire to continue to speak as if one were a *bona fide* Moorean’.¹⁰⁸ The challenge thus seems to remain: has the subjectivist earned the philosophical right to speak in such metaphysical-sounding terms as *supervenience* and *constitution*? I do not claim to have resolved this issue. It is probably correct that value is not going to be to the subjectivist all that it is to the objectivist, but perhaps this is because the objectivist is asking for it to be too much. All that I claim at this point is that the subjectivist has an equal claim to that of the objectivist to try to analyse such notions as intrinsic value. At this point I do not claim that subjectivist analyses are preferable over objectivist ones; this would require further argument and ultimately lies beyond the scope of this work.

¹⁰⁶ Gauthier (1986) (p. 47)

¹⁰⁷ Rabinowicz & Österberg (1996) (p. 19). Rønnow-Rasmussen (2003) suggests that though many subjectivists probably are projectivists, the subjectivist is not necessarily committed to projectivism. He writes ‘Establishing a relation between an object and a final attitude is what is needed (perhaps all that is needed) to generate an act of constitution’ (p. 253). This issue will be picked up on again in Chapter 6, where I will discuss value projection and the ontological status of value.

¹⁰⁸ Brännmark (2002) (p. 66)

VII. First-Person and Third-Person Perspectives

Another way to formulate the difference between subjectivist and objectivist views of value is in terms of the first- and third-person perspectives. The first-person perspective can be understood as the perspective from the inside, looking out, whilst the third-person perspective can be understood as the perspective from the outside, looking in. The latter can perhaps be assimilated with Henry Sidgwick's 'point of view of the universe'¹⁰⁹, or Thomas Nagel's 'view from nowhere'.¹¹⁰ Subjectivists hold that it is only from the first-person perspective that things can be good or bad, better or worse, etc. Things matter to particular people, but nothing matters from the point of view of the universe. This conclusion flows quite naturally from the idea that value is grounded on attitudes, since the universe does not seem to have any attitudes.¹¹¹

Objectivists, on the other hand, think that they can see values from the third person perspective. Jan Österberg, the champion of the satisfaction interpretation of preferentialism, notes correctly that the first-person perspective (which he calls the immanent perspective) does not give us access to objective values. Instead, he holds, the first-person perspective 'gives rise to a systematic misapprehension: it makes us believe that the values we acknowledge in that perspective exist independently of us, whereas in fact they are projections of our preferences.'¹¹² But when we view things from the third-person perspective (which he calls the transcendental perspective—the perspective of a detached spectator) we can, Österberg insists, see that preference satisfaction is finally valuable and that nothing but preference

¹⁰⁹ Sidgwick (1981). The phrase occurs in the following passage: 'by considering the relation of the integrant parts to the whole and to each other, I obtain the self-evident principle that the good of any one individual is of no more importance, from the point of view (if I may say so) of the Universe, than the good of any other' (p. 382). (Of course, if nothing is good or bad from this cosmic perspective it is trivially true that nothing is more important from this perspective than anything else.)

¹¹⁰ Nagel (1986)

¹¹¹ There are, of course, attitudes *in* the universe, but that is another matter entirely.

¹¹² Rabinowicz & Österberg (1996) (p. 12)

satisfaction is thus valuable. Other objectivists would, of course, insist that we can see from the same perspective that other things are finally valuable.

Subjectivists hold that we can see nothing of the sort. As Rabinowicz writes,

From such a perspective [the transcendental, third-person perspective of a detached spectator] values are no longer to be seen. The world is drained of evaluative meaning. In order to save values, we have to retain the immanent [first-person] perspective.¹¹³

Simon Blackburn voices the same idea, applied to ‘moralizing’ rather than valuing in general, when he says, about the third-person perspective (or as he calls it, ‘this sideways theoretical perspective’),

It is not that this perspective is illegitimate, but it is not the one to adopt for finding ethical truth. It would be if such truth were natural truth, or consisted of the existence of states of affairs in the real world. That is the world seen from the viewpoint that sees different and conflicting moral systems—but inevitably sees no truth in just one of them. To ‘see’ the truth that wanton cruelty is wrong demands moralizing, stepping back into the boat, or putting back the lens of a sensibility.¹¹⁴

Some may perhaps want to go further than Blackburn and declare the third-person perspective to be somehow illegitimate or impossible as such. What is important here, however, is that all subjectivists hold that seeing something as good or bad requires seeing the object in question through the lens of one’s pro and con attitudes. If this lens is removed—as it is when we move to the third-person viewpoint—good and bad are simply lost from sight. One must step back into the boat and put the lens of a sensibility back in order to see goodness and badness.¹¹⁵ This idea goes well together with the

¹¹³ Ibid (p. 19)

¹¹⁴ Blackburn (1988) (p. 372). Blackburn speaks here of ‘truth’. I leave it open at this stage whether subjectivists can or would want to recognize evaluative truths. I return to this issue in Chapter 8.

¹¹⁵ An objectivist could here agree that desires are necessary in order for us to perceive values, but deny that values are products of our desires. It seems to me that there is an obvious inherent tension in such a view; for what reason could there be to think that values exist

previously recognized one that attitudes are always present in the background of valuing.

On one interpretation of David Hume, he implicitly appealed to the distinction between first- and third-person viewpoints in this oft-quoted passage from *Treatise*:

Take any action allow'd to be vicious: Wilful 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 never can find it, till you turn your reflexion into your own breast, and find a sentiment of disapprobation, which arises in you, toward this action. Here is a matter of fact; but 'tis the object of feeling, not of reason. It lies in yourself, not in the object.¹¹⁶

As long as you examine the situation from the outside, from an objective, third-person perspective, you will find 'only certain passions, motives, volitions and thoughts'. No matter how closely you look, values cannot be seen from that perspective. It is first when you return to the first-person perspective—when you 'turn your reflexion into your own breast' and thus come to view the situation from within—that you will be able to see the badness of it. If values were objective it would presumably have been possible to see them from the third-person perspective, but because the ground of value 'lies in yourself, not in the object', this is not possible.¹¹⁷

independently of us, if they are accessible only from within? However, I do not attempt to resolve that issue here.

¹¹⁶ *Treatise* (iii.i.i.)

¹¹⁷ On Martin Harvey's (2004) interpretation of Thomas Hobbes, Hobbes too operated with first and third person perspectives. According to Harvey, from the third-person point of view, Hobbes presents 'an objective "civil science", which identifies the universal "rules" for securing peace and avoiding war'. The challenge for Hobbes involves ensuring that people actually want follow these rules. Harvey contends that 'Hobbes's axiology, though formally subjective, places no substantive restrictions on the type of value that individuals may place upon justice' (pp. 439–440). 'Axiological room obtains for construing justice as either

VIII. A Bleak View on which Nothing Matters?

Derek Parfit's attack on subjectivist views rests fundamentally on his insistence that things matter from an impersonal point of view. Subjectivism, he holds, is a 'bleak view on which nothing matters'. He notes that 'some subjectivists would admit that, on their view, nothing matters in an impersonal sense. It is enough these writers claim, that some things matter to particular people'. This is indeed an intuition that I believe subjectivists feel as strongly as Parfit feels his intuition. But insisting on an intuition is unlikely to convince anyone, and Parfit is right when he says that 'this reply shows how deep the difference is between the two kinds of theory'.¹¹⁸ For a subjectivist it is hard to understand how anything could possibly matter beyond mattering to particular people. The cosmos just doesn't care.

Objectivists will of course continue to insist that they can see values from the third-person perspective, but according to the subjectivist, objectivists thereby illegitimately import something to the third-person perspective that essentially belongs to the first-person perspective. Values that are actually subjective are mistakenly thought to be objective. This may be a natural mistake, but a mistake it is nonetheless. As previously noted, subjectivists need not deny that values seem to have objective validity. Indeed, this would help to explain the persistent attraction of objectivism. But according to the subjectivist this seeming objectivity of value is an illusion.¹¹⁹

instrumentally or intrinsically valuable (relative to the desires of the individual agent)' (p. 445). The rules of justice can thus be said to have value, and to be normative, but only from within the first-person perspective of the individual agent.

¹¹⁸ Derek Parfit (2011) (p. 107)

¹¹⁹ In Chapter 8 I discuss whether this diagnosis of systematic error necessarily leads the value subjectivist to an error theory. I argue there that while an error theory is one possible option for subjectivists to take, she is not necessarily committed to such a theory. All value subjectivists hold that insofar as people in general hold objectivist second order beliefs about value they are systematically mistaken about the nature of value. But this, by itself, amounts to no more than holding that value objectivism is simultaneously widespread and false. This is not enough to make one an error theorist. In order to arrive at an error theory one must in

On the subjectivist picture, what can be seen from the third-person perspective are only our various attitudes—only ‘certain passions, motives, volitions and thoughts’, as Hume put it. We can infer from what we see from the third-person viewpoint that value constitution is going on whenever subjects hold pro and con attitudes. We can further see what the objects of their respective attitudes are and the non-evaluative properties of these objects in virtue of which they are favoured or disfavoured. But whilst this information allows us to conclude that an object *O* is good *relative to* subject *S*, it does not allow us to conclude that *O* is *good* or that *O* is *good for S*. Statements of the good-relative-to kind are not themselves evaluative or normative, and that many statements of this kind are objectively true affords us no more ground for thinking that we have detected an objective value than does *O* is good *according to S*. Trying to derive genuinely evaluative or normative conclusions strictly from what is given in the third-person perspective would be like trying to squeeze blood from a stone.

addition accept certain specific claims about the nature of evaluative language, claims that, as we will see, not all subjectivists accept.

3. THE IMPORTANCE OF BEING REAL: VALUE AND HYPOTHETICAL ATTITUDES

[A]ny variation of interest or of its object will determine a variety of value, . . . any derivative of interest or its object will determine value in a derived sense . . . [and] any condition of interest or its object will determine a conditional value. In short, interest being constitutive of value in the basic sense, theory of value will take this as its point of departure and centre of reference; and will classify and systematize values in terms of the different forms which interests and their objects may be found to assume.

Ralph Barton Perry, *General Theory of Value*

I. Introduction

A common, and indeed tempting, way to capture the relation between value and attitudes is to say that value depends on attitudes in such a way that if there were no attitudes, there would be no values, and if people held different attitudes than those that they actually have, different things would be valuable. It is easy enough to imagine situations in which people have very different desires and preferences from those that we in fact have. Take a paradigmatic example of a worthless thing such as a saucer of mud, and equally paradigmatic examples of good things such as knowledge, peace and biodiversity. You can imagine a hypothetical situation in which the saucer of mud is desired for its own sake, and where people lack any pro attitudes

towards knowledge, peace or biodiversity. Is the subjectivist forced to conclude that the saucer of mud would be finally valuable in that situation, whilst knowledge, peace and biodiversity would lack value?

In the present chapter I consider some objections against value subjectivism that can be understood as asking us to imagine situations in which we hold different attitudes than those that we in fact hold. Insofar as our attitudes would have been different from what they are, it seems that what is good and bad would also be different. I argue that this is not so, however, because it is always the attitudes that we actually have that ground values even in hypothetical situations.

II. What If There Were No Subjects with Attitudes?

Imagine first a world full of natural beauty that is not inhabited by any beings with affective states. Since there are no attitudes in such a world, it would seem that a subjectivist is committed to saying that there can be no value in such a world. But consider again a final desire for the survival of the rainforest. Insofar as such a desire extends to the hypothetical worlds in which there exist no subjects with attitudes—that is, insofar as we wish, for its own sake, that the forest would be here in all its glory even if there were no creatures like us around to appreciate it—then subjectivism implies that the survival of the rainforest would retain its final value even in such worlds. This is so, as we have seen, since on this view value is taken to supervene not on the attitudes involved but on the objects of the attitudes, in virtue of the non-evaluative features for which the objects in question are being favoured.

The objectivistic satisfaction interpretation, by contrast, would imply that the survival of the rainforest would be of no value in worlds without entities with affective states, as it would not satisfy any preferences there. Since, on the satisfaction view, value is taken to supervene solely on states that involve the satisfaction of preferences, in worlds in which there are no subjects with preferences, there would be nothing for value to supervene on and hence no value in such worlds.

Adam D. Moore raises the following objection:

If value is intimately tied to our affections and value is brought into the world only when a desire or preference is satisfied, then entities that do not have affective states are left out of the moral picture in terms of value. Subjectivist accounts of value would thus be elitist in the sense that only some living entities would be able to produce value by having intrinsic desires satisfied.¹²⁰

Armed with the distinction between the satisfaction interpretation and the object interpretation, the subjectivist's response to this objection should by now be predictable: one doesn't need to be able to produce value in order to be a bearer of value. It is indeed true that, on subjectivist views, only some living entities are able to produce value (though not, as we have seen, by having their desires satisfied, but simply by having desires or other pro and con attitudes). But it is not true that entities that do not have affective states are thereby 'left out of the moral picture in terms of value'. To the contrary, such entities can and do enter that very picture as objects of our attitudes. And it is the objects of our attitudes that are taken to carry value on the object interpretation, which I have identified with subjectivism.

On the subjectivist view, it is only by entering as objects of our attitudes that things have value, but this does not imply that entities that themselves lack affective states therefore lack value. A multitude of such entities are, as a matter of fact, desired (admired, loved, respected, etc.) for their own sakes. And since many entities beyond those that themselves have affective states actually figure among the objects of our final attitudes, subjectivism does not have the implication that such entities lack final value. Even in worlds in which there are no entities with affective states, and thus no producers of value, things can still be finally valuable in those worlds. The values of those hypothetical objects are constituted by and grounded on our actual attitudes. It might, of course, be insisted that granting things like rainforests the status of value bearers is not sufficient and that subjectivism is faulty because it denies such things the status of value producers. But this would seem to beg the question against the subjectivist.

¹²⁰ Moore (2004) (p. 80)

III. What If We Had ‘Bad’ Attitudes?

If value is tied to attitudes, it would seem that if we had different attitudes, different things would be valuable. We can imagine a hypothetical world in which people have very different attitudes from those that they actually have. To reuse my earlier example, we can imagine that people lack any pro attitude towards paradigmatically good things such as knowledge, peace and biodiversity. As argued above, the idea that value is determined by attitudes is open to two radically different interpretations. Depending on which interpretation one adopts, one would address this concern in very different ways.

The objectivistic satisfaction interpretation adjusts itself to those hypothetical attitudes that we would have *in* these hypothetical cases. The defender of the satisfaction view thus must conclude that in a situation in which no one cared for knowledge, peace or biodiversity, these things would lack value, since they would not satisfy any preferences in that situation!¹²¹ The object interpretation follows instead our actual attitudes *for* hypothetical cases. This means that insofar as we actually do have final preferences for things such as knowledge, peace, and so on, that extend to hypothetical cases in which no one cares for such things, the object interpretation implies that these things are finally valuable even in those hypothetical situations in which these things are not favoured. Correspondingly, on this interpretation, a saucer of mud would lack value even in those hypothetical situations in which it is desired for its own sake. It is the attitudes of actual subjects that determine the value of both actual and hypothetical objects.

But this does not mean, however, that the preferences we could have had but don't actually have are unconditionally disregarded by the object theoretician. Rather, such hypothetical preferences enter into the picture insofar as our actual preferences are sensitive to the hypothetical preferences in question. Whilst many of our actual preferences are not thus sensitive,

¹²¹ It should be kept in mind, though, that unlike the object interpretation, the satisfaction interpretation does not ascribe final value to such things even in the actual situation in which these things *are* desired for their own sake; on the satisfaction view, as we have seen, only preference satisfaction carries final value.

some clearly are. To reuse an earlier example: When I imagine a situation in which I don't desire ice cream, I fail to see any value in me having ice cream in the imagined situation. In other words, my desire for ice cream is an example of a desire that does not extend to the hypothetical case in which I don't desire ice cream. It is a case where my actual desire for a hypothetical situation is sensitive to my hypothetical (lack of) desire. But when I imagine situations in which I do not care about knowledge or biodiversity, for example, my reactions are very different. I then emphatically disregard my hypothetical (lack of) pro attitudes. In both types of case the values involved are determined by my actual attitudes.

Another way to put this is in terms of the by-now familiar distinction between ground and object. Just as an attitude and the object of that attitude need not be located in the same place (I can love someone who is in a different country) and time (I can admire Socrates even though he is long dead), so too can an attitude and its object be located in different possible worlds (I can here-and-now wish that were I outdoors where it is raining I would have an umbrella with me). In other words, a favoured object—and thus the non-evaluative properties of that object in virtue of which it is being favoured—may be located in a different possible world than the attitude of which it is an object. On subjectivist views, the constitutive ground of a value is identified with an actual attitude, but the bearer of that value—the object of the attitude—need not itself be actual. All values are grounded in actual attitudes, but the objects of value, and thus the non-evaluative properties of these objects upon which their value supervenes, need not all be located in the actual world. Our actual attitudes are thus taken to constitute value in this world as well as in other possible worlds. The constitutive power of our attitudes does not respect the boundaries between the actual and the merely possible, any more than it respects the boundaries between past and present or between near and far. Indeed, anything that can be an object of a pro or con attitude can be an object of value, irrespective of the object's location in space, time, or modality.

Yet another way to understand this is by remembering the projectivist nature of subjectivism. On such views, as we have seen, value is not something that forms part of the mind-independent world but rather is something that we project upon the world, or more precisely, as Rabinowicz points out, 'upon the whole set of possible worlds'. Value is thus projected not only onto objects in the actual world but upon objects in other possible worlds as well. In this sense, 'while not world-bound in [its] *range of application*' value is 'an expression of a particular world-bound *perspective*.

the perspective determined by the preferences we actually have.¹²² In this sense value is invariant over possible worlds, whilst always remaining relative to a particular world-bound perspective.

Perhaps the following analogy can shed more light on this aspect of subjectivism. Consider our usage of the expression *one meter*. This particular length is conventionally fixed by a relation to some standard, such as the standard meter in Paris.¹²³ Consider an object with a certain length, such as the author of this book: 1.78 meters. Consider another possible world in which the standard was different. We would still say that I am 1.78 meters tall in that world, even if in that world, I would correctly say that I was, say, 3.5 meters tall. In Saul Kripke's term *one meter* is a so-called 'rigid designator'.¹²⁴ On the subjectivist view, ascriptions of value might be seen to function similarly: actual preferences determine which objects or states of affairs have value, and these objects remain valuable in those worlds in which people have different preferences.

Is the Restriction to Actual Attitudes Ad Hoc?

The objectivist Michael Huemer acknowledges that one way to counter the objection that tying values to attitudes leads to counter-intuitive results when applied to hypothetical situations in which we have different attitudes than we in fact have

would be to 'rigidify' evaluative terms—that is, one might propose that 'good' always has for its extension what we actually approve of, even when we are discussing counter-factual situations in which we wouldn't approve of those things. Thus, the reason we don't have to say that if everybody approved of Hitler then Hitler would be good, is that we presently

¹²² Rabinowicz & Österberg (1996) (p. 19)

¹²³ A similar example can be found in Rabinowicz (1996) (p. 12), where it is ascribed to Krister Bykvist.

¹²⁴ Kripke (1980)

disapprove of Hitler, and we continue to disapprove of him even when we're thinking about that counter-factual situation.¹²⁵

Now, it should be noted that this is not quite what I take the subjectivist to be saying.¹²⁶ Huemer is talking about evaluative terms (and their extension) rather than about value as such. Nonetheless, what Huemer goes on to say may be taken to target also the subjectivist. He calls this focus on actual attitudes an *ad hoc* manoeuvre.¹²⁷ But is it really *ad hoc* to restrict the attitudes that are taken to ground value to actual attitudes? After all, given the generic subjectivist idea that value is grounded on attitudes, it should hardly be surprising that it is attitudes that are actually held by someone that are taken to play this value-grounding role rather than some attitudes that nobody actually has. It therefore seems quite natural, within a subjectivist framework, to hold that it is on our actual attitudes that values are grounded.

Furthermore, the conclusion that it is our actual attitudes that determine the perspective from which values are constituted flows quite naturally from the idea that value-constitution requires the first-person perspective; the idea that it is only from the first-person perspective that things are good or bad, better or worse, etc. Following Simon Blackburn, I previously likened the activity of valuing something with seeing it through the lens of a sensibility. This lens is here identified with the present pro and con attitudes of actual subjects. Whether the things that are being objects of our pro and con attitudes are actual or merely possible (or whether they are

¹²⁵ Huemer (1996) (Section IV)

¹²⁶ David Lewis (1989) says more or less what Huemer describes: 'We might have been disposed to value seasickness and petty sleaze, and yet we might have been no different in how we used the word "value". The reference of "our actual dispositions" would have been fixed on different dispositions, of course, but our way of fixing the reference would have been no different. In one good sense—though not the only sense—we would have meant by "value" just what we actually do. And it would have been true for us to say "seasickness and petty sleaze are values" (pp. 132–133). I take it that the problem Lewis is here discussing arises for his dispositional theory of value in virtue of it 'advancing an analytic definition of value' (in terms of what we [actually] desire to desire'). But this problem does not arise for subjectivism (as I understand that view here), at least not in the same way, because my subjectivist offer no definitions of evaluative terms.

¹²⁷ Huemer (1996) (Section IV)

in the present or belonging to a different time), does not alter the nature of value or its ground. Hence, whether we are favouring something actual or something hypothetical (or something in the past or in the future), the constitutive ground of the value is always identified with our actual (and present) attitudes.

A Dictatorship of the Actual?

Is the actualist character of value subjectivism objectionable? Krister Bykvist has called it a form of ‘dictatorship of the actual’.¹²⁸ Could anything further be said in order to dispel or lessen qualms about the privileged status that subjectivists give to actual over merely hypothetical attitudes? As Rabinowicz points out, it is difficult to provide a justification for this bias for the actual over the merely possible simply because the bias in question is so fundamental for us humans. The line between the real and the unreal seems to be of utmost importance, even if its importance cannot be justified in terms of anything else. What the subjectivist can do is simply to remind the potential critic of the importance of being real.¹²⁹

IV. What If Our Attitudes Are Not as We

Suppose?

As we have seen, subjectivism does not—contrary to what is commonly assumed—have the implication that if people had different attitudes than those that they in fact have, different things would be valuable. However, subjectivism does imply that insofar as people’s attitudes are *in fact* different

¹²⁸ Bykvist (1996) (p. 6)

¹²⁹ Rabinowicz & Österberg (1996) (p. 23)

from what we suppose them to be, then what is in fact valuable (relative to actual people) is different from what we suppose it to be.

One way in which we might think about this would be to let another possible world function as the actual world. We would then get a different distribution of values. But with every choice of actual perspective, value will be invariant across possible worlds. An object that is valuable given a certain perspective thus remains valuable relative to this perspective, even in other possible worlds in which our attitudes toward it are different. There is thus a relativist element built into the value subjectivist view, but it is a relativism that is compatible with value being invariant across possible worlds relative to a fixed perspective.¹³⁰

Michael Huemer finds this problematic. ‘One consequence that can not be escaped by this *ad hoc* manoeuvre’ to rigidify evaluative terms ‘is that any sentence of the form, “If we approve of x, then x is good”, where the conditional is indicative (not subjunctive), must be true.’¹³¹ Even if, as Huemer concedes, subjectivists are not committed to affirm counter-factual claims such as ‘if most people *had* approved of starvation, then starvation *would* have been good’, he thinks that they must affirm conditional statements like ‘if most people *do* approve of starvation, then starvation *is* good’ (Huemer’s examples). Accepting the latter conditional might be unacceptable even if it is acknowledged that the antecedent is plainly false; it is clearly not the case that most people approve of starvation. But whilst the position Huemer has in mind when he poses this objection will affirm such evaluative truths, this is not true of subjectivists as I understand them here. ‘If most people approve of starvation, then starvation is good’ seems to me to be a first order, evaluative statement, and one that should be rejected by subjectivists and objectivists alike. Only non-evaluative claims to the effect that ‘if most people approve of starvation, then starvation is good *relative to these people*’ would come out as true on subjectivism.¹³² And affirming this

¹³⁰ Though all forms of subjectivism are relativist in this sense, in Chapter 5 I will use the term *relativism* in a different way that allows me to distinguish between relativist and absolutist versions of subjectivism.

¹³¹ Huemer (1996) (Section IV)

¹³² It should be added also that, as I will argue more fully in the next chapter, something is good relative to someone only when that someone has an accurate picture of the thing in question. If the favoured object in fact lacks some or all of the properties for which it is being favoured by the relevant subject, then the subject has failed to confer value on that object

claim does not commit one to approve of starvation or to hold it good, or to hold it good *for* anyone.

Still, even these implications might perhaps offend some sensibilities. David Lewis remarks about his own view of value that even if it does not ‘imply that if we had been differently disposed, different things would have been values . . . it comes too close for comfort’¹³³ But I doubt that this remaining discomfort can form the basis of an independent objection against subjectivism without presupposing the falsity of the very idea that values are grounded on attitudes.

V. Invariance and Consistency

Invariance

As noted above, if we let another possible world function as the actual world, we could get a different distribution of values. But with every choice of actual perspective, value will be invariant across possible worlds. An object that is valuable given a certain perspective thus remains valuable relative to this perspective even in other possible worlds in which our attitudes towards it are different. Rønnow-Rasmussen discusses what he calls the invariance thesis, which says ‘The final value of an object is invariant over possible worlds’. It follows from this thesis that ‘if an object is valuable, the object (or any object exactly similar to it) carries this value in whatever

relative to him or herself. With this in mind, the subjectivist is not committed even to ‘If most people approve of starvation, then starvation is good relative to these people’ but only to ‘If most people approve of starvation *whilst having an accurate picture of what it involves*, then starvation is good relative to these people’.

¹³³ Lewis (1989) (p. 113)

possible world in which it is present'.¹³⁴ As Rønnow-Rasmussen notes, it is a standard objection to subjectivist accounts that they have to reject the invariance thesis.

[I]f we fix value to desires and other kinds of attitudes then an object, *x*, would not be valuable (i.e., its value would not have been constituted) in a possible world in which *x* is not the object of a desire. If the world were different, with a different set of desires, there would be another set of valuable objects than what is actually the case.¹³⁵

We are now in a position to see that this objection rests on a mistake. As argued above, whilst the constitutive ground of value is identified with an actual attitude, the object upon which the value is taken to supervene need not itself be located in the actual world. Given my actual attitude towards *x*, the value of *x* will be invariant over possible worlds. Subjectivists can thus accept the invariance thesis. However, as Rønnow-Rasmussen points out, unqualified the invariance thesis is unacceptable for reasons that are independent of subjectivism. If the value of an object supervenes on its non-essential properties, then—as both objectivists and subjectivists should agree—in a world in which the object lacks some of these properties, it will no longer be valuable. Rønnow-Rasmussen's example of this is a valuable painting that would not be valuable had it been painted with different colours. Assuming that the colour properties are not essential properties of the painting, subjectivists and objectivists should agree that the painting need not retain its final value in worlds in which it has different colour properties. My previous example of a guitar that is valuable in virtue of having once belonged to Steve Howe may be used to illustrate the same point. Having belonged to Howe is not an essential property of the guitar (had Howe never owned it, it would still be the same guitar), and in those worlds in which Howe never owned the guitar, it would not have the same value. The difference between the two examples is that the colours of the painting are internal (yet non-essential) properties of it, whilst having once belonged to Steve Howe is a relational property of the guitar. In both examples, the respective features of the objects upon which their values

¹³⁴ Rønnow-Rasmussen (2003) (p. 257)

¹³⁵ *Ibid*

supervene are themselves altered, but it is still the case that in all of the possible worlds in which the supervenience bases remain the same the objects retain their values. Whether final value may supervene on non-essential properties depends on the nature of value bearers. If the only bearers of final value are states of affairs, then all valuable objects will have their values in virtue of essential properties because abstract entities like states of affairs have all of their internal properties essentially; if one of the features of a state of affairs was different it would be a different state of affairs. This complication will therefore not arise for those subjectivists and objectivists who accept a form of value-bearer monism according to which only states of affairs can be bearers of final value.

Rønnow-Rasmussen uses the invariance thesis to distinguish four different variants of value subjectivism. According to what he calls first-person perspective subjectivism, the invariance thesis is true given the actual first-person perspective but false relative to the third-person perspective.

Thus, given my actual perspective of what I here and now prefer, the value of whatever I value will be invariant across worlds. But as soon as I step out of this first person suit, and look upon the whole situation from the position of a detached observer, it will be clear that value may in fact vary depending on what preferences there are. Or to be more accurate, value is still dependent on the subject's 'actual' preferences; it is just that *what the actual perspective is may vary depending on from which world the preference stems*.¹³⁶

On this view, as I interpret it, from the actual first-person perspective, actual attitudes enjoy a privileged status over merely possible ones. But from the third-person perspective no possible world is privileged over any other. If we would allow that every possible world is actual from its own perspective, we must allow that—insofar as people in other possible worlds have attitudes that extend to our world—people in these worlds constitute values in our world just as we do in theirs. Given the actual first-person perspective value is invariant across possible worlds, but from a third-person perspective we can speak only of things being good or bad relative to this or that (world-bound) perspective. As previously noted, whilst the constitutive ground of value is tied to a particular world-bound perspective—the perspective

¹³⁶ Ibid (pp. 258–259)

determined by the attitudes that we actually have—values are not world-bound in their range of application. The values constituted by actual subjects may thus well apply not only to us in the actual world, but to all possible subjects (regardless of whether they share our attitudes or not). From our actual perspective, the values constituted by our attitudes may apply to every possible subject. And even though the values constituted by merely possible subjects' attitudes may likewise apply to us, they do so only from their perspective. These values are thus only values relative to such possible subjects.

Exclusive first-person perspective subjectivism also affirms that the invariance thesis is true given the actual first-person perspective, but it remains silent on the third-person perspective, perhaps because the very idea of such a detached perspective is taken not to make sense.

On proper actualism, the invariance thesis is true given the third-person perspective as well as from the first-person perspective. Unlike first-person perspective subjectivism, proper actualism recognizes only one set of value-grounding attitudes. The only attitudes that have the power to constitute values are those of the inhabitants of the actual world. Merely possible subjects' attitudes are taken to be impotent in this respect. As Rønnow-Rasmussen points out, we might question to what extent proper actualism really is a version of subjectivism, especially if subjectivism is taken to entail some version of value relativism. He chooses to regard proper actualism as 'something in between subjectivism and objectivism', a kind of hybrid view with both subjectivist and objectivist elements. The view is subjectivist in holding that values are constituted by subjects and that value supervenes on the features of objects for which they are favoured, but it is supposedly objectivist in allowing us to speak of things as being good or bad, not only from the first-person perspective, but also from the third-person perspective. However, I do not believe that proper actualism must be seen to have this objectivistic implication. Proper actualism may be stripped of its objectivist element and thus be made fully compatible with subjectivism as I have described it. Interpreted in this way, the proper actualist, like the first-person perspective subjectivist, holds that what can be seen from the third-person perspective is only that things are good or bad relative to perspectives. But, unlike the first-person perspective subjectivist, the proper actualist holds that things are good or bad only relative to actual subjects. What is good or bad relative to non-actual perspectives is either considered a moot question or simply answered by saying that nothing is either good or bad relative to non-actual subjects.

Finally, Rønnow-Rasmussen adds variance subjectivism, which rejects the invariance thesis wholesale.¹³⁷ Unfortunately he does not tell us on what grounds the variance subjectivist holds that the invariance thesis is false given the first-person perspective. It is therefore hard to evaluate the credentials of this view as a version of value subjectivism. In previous sections I have spoken as if I am following first-person perspective subjectivism, and I will continue to so speak in what follows. However, I do think that both exclusive first-person perspective subjectivism and my streamlined version of proper actualism remain possible alternatives for the subjectivist.

Consistency

When asked why you value a particular thing but not another, you are expected to be able to point to some difference between the two objects that goes beyond one being good and the other bad. It is commonly assumed that if an object—let's call it O_1 —is taken to be good, and a different object, O_2 , is taken to be not good, then this must be because of some difference in the respective non-evaluative properties of O_1 and O_2 . If we agreed that there is no non-evaluative difference between O_1 and O_2 , it must also be agreed that there is no evaluative difference between them. This is a feature of supervenience, which I argued in Chapter 1 to be common ground among subjectivists and objectivists, something they both need to respect and explain.

Subjectivists have sometimes argued that this phenomenon is readily accounted for given that value is constituted by and grounded on attitudes. Indeed, some subjectivists have argued that they can give a better explanation for this than can objectivists. Simon Blackburn, for example, says that

it is not possible to hold an attitude to a thing because of its possessing certain properties and, at the same time, not hold that attitude to another thing that is believed to have the same properties. The nonexistence of the

¹³⁷ Ibid (p. 260)

attitude in the second case shows that it is not because of the shared properties that I hold it in the first place.¹³⁸

The idea seems to be that insofar as I have a certain attitude towards O_1 in virtue of O_1 having non-evaluative properties P_1, \dots, P_n , then it is impossible for me not to have this same attitude also towards O_2 which (I believe) also has P_1, \dots, P_n . If it turns out that I nonetheless lack the attitude towards O_2 , this must be because it is not in virtue of P_1, \dots, P_n that I have the attitude towards O_1 after all, but rather in virtue of some other property that O_2 lacks. This makes sense given that our attitudes are responses to (and presumably caused by recognition of) the non-evaluative properties of the objects in virtue of which we favour or disfavour them. And if this is so, and if values are taken to be grounded on attitudes, it follows that objects (that are thought to be) exactly alike in their non-evaluative features cannot (be thought to) have different values.

Still, it might well be insisted that it is, contra Blackburn, quite possible for an individual subject to have an attitude towards O_1 but at the same time to lack a similar attitude towards O_2 , even though the two objects are (believed to be) exactly alike in all their non-evaluative features. I will discuss several different ways in which a subjectivist might go about dispelling this worry. But first we need to know more about the particular case at hand, including what kind of objects we are dealing with: Are they, for example, concrete objects like paintings or persons, or abstract objects like states of affairs? What kind of values are involved, is it intrinsic or extrinsic value? Must the objects in question be objects of subjects' conscious attention in order for an act of value constitution to take place?

First, if we assume that O_1 and O_2 are numerically distinct, but qualitatively identical concrete objects, is it then possible for someone to have an attitude towards O_1 without thereby having the same attitude towards O_2 ? If O_1 and O_2 are exactly alike only in their internal features, they might well differ in their relational features, and this might make all the difference in our attitudes. If, to reuse an earlier example, two guitars have exactly the same internal properties, the one may still be favoured over the other in virtue of standing in the relation of having belonged to a particular

¹³⁸ Blackburn (1971) (p. 122)

person. The two guitars then differ in their respective extrinsic value. But in this case consistency does not break down since the non-evaluative features on which the respective values are taken to supervene are different.

But what if the objects are exactly alike in all their non-evaluative features, internal as well as relational? If one knows this to be so, can one still withhold an attitude from O_2 that one has towards O_1 ? This is certainly unusual, and perhaps it is not even factually possible. The fact that our attitudes normally don't and perhaps never behave in such highly uneven ways may, I believe, be a large part of the explanation for why we have the consistency intuition in the first place. This is why we find it so startling that anyone would dole out his or her favour or disfavour in such a lopsided manner. This, in turn, can perhaps explain why the consistency requirement is a conceptual necessity: the consistency intuition has come to be incorporated into conventional evaluative language. Words like *good* and *bad* are governed by a special logic that forbids applying them in ways that violate this consistency. This explains why competent language users would find it odd were someone to say 'This and that are exactly alike, except the one is good and the other is bad'.

Still, however unlikely it is to obtain in the real world, it could be insisted that it is logically possible for a subject to withhold an attitude from O_2 even though it is thought to share the properties P_1, \dots, P_n for which O_1 is favoured, and that this possibility remains a problem for the subjectivist who wants to respect the consistency requirement. It may be asked how the subjectivist can guarantee that consistency holds in every conceivable circumstance. A bad answer to this question would be to say that the object of the subject's attitude is really just the set of properties P_1, \dots, P_n . Though this would solve the problem at hand, that objects are favoured *for* some of their properties does not imply that it is really those properties that are being favoured rather than the objects that are thought to possess the properties. Besides, this would imply an unjustified restriction on the type of objects that could be bearers of value (only sets of properties).

Another obvious solution would be to impose a restriction on the attitudes that ground value. Attitudes that would otherwise ground values but that violate consistency are disqualified from doing so and thereby robbed of their value-constituting power. In this way a guarantee that consistency always holds would be built into the theory from the start. The rationale for such a restriction could appeal to the consistency intuition or to the conceptual necessity of the consistency requirement that is thought to be common to all involved in the debate (being a feature of supervenience).

Still, one might regard this as a somewhat *ad hoc* move. There may, of course, be independent reasons for imposing some requirement of this kind. Idealization and restrictions of value-grounding attitudes will be discussed further in the next chapter.

A better solution has been suggested to me by David Alm. On this proposal, what is going on when a subject is having an attitude towards something is only loosely described by saying that value is being conferred upon an object (the value bearer). Strictly speaking, it is rather the status of being value-makers that is being conferred onto the properties of the object for which it is being favoured or disfavoured. Hence, by favouring or disfavoured an object in virtue of some of its properties, the subject bestows upon these properties the status of being value-makers relative to the subject. If, for example, a subject favours O_1 on account of properties P_1, \dots, P_n , and this is taken to bestow upon P_1, \dots, P_n the status of being value-makers relative to the subject, it follows that O_2 (as well as any other object) that also possesses P_1, \dots, P_n is *ipso facto* also valuable relative to the subject in question.

But what if the subject favours O_1 on account of P and disfavors O_2 on account of P ? Has she now conferred upon P the status of being a value-maker and at the same time the status of being a disvalue-maker? We seem once more to have a pernicious inconsistency on our hands. But perhaps this has all been a bit too quick. We may not want to say that a property that is a value-maker in one context necessarily plays this role in every context. Hence, one response to the present worry would be to say that one and the same property can be a value-maker in one context and at the same time not be a value-maker in another context. Whether a particular property is a value-maker or not in a particular context may depend on the presence or absence of certain other features. The same may hold for a set of properties: that P_1, \dots, P_n make a particular object good in one context need not imply that P_1, \dots, P_n make any object that has them good in every context. A subjectivist may want to follow Dancy in saying that features that are not themselves value-makers (such as enabling and disabling conditions) might make a difference to the value of an object. A subject favouring O_1 on account of P_1, \dots, P_n and disfavouring O_2 on account of P_1, \dots, P_n , need therefore not lead to a breakdown of consistency given the presence or absence of certain other relevant features. Dancy makes the point in terms of wrongness rather than goodness:

It is extremely plausible to suggest that if an action is wrong, every other action that is exactly similar in non-moral respects must be wrong also. . . . It is nothing like so plausible to suggest that if an action is wrong, every other action that shares the features that make the first one wrong must also be wrong. Two actions may be similar to each other in a limited way, that is, in the respects that disfavour the first one and thereby make it wrong, but differ in other respects so that the second is not wrong; the features that manage to make the first one wrong are prevented from doing so in the second case because of variations that lie beyond the common [features].¹³⁹

Dancy takes this point to be established by appeal to the distinction between favourers and enablers that I first introduced in Chapter 1.

But if we include in P_1, \dots, P_n not only the value-making features of the object (internal as well as relational) but all the features that make a difference to the value of the object (the complete supervenience base, including enablers and disablers, etc.) then we may once more ask whether it is possible for a subject to take differing attitudes towards two objects that share all of these features. Insofar as the answer is yes, it seems that subjectivism does lead to a breakdown of consistency. Here it seems that the subjectivist may need to either fall back on Blackburn's strategy (though taking into account that there might be more to supervenience than value-making features), thus declaring attitudes that lead to a breakdown of consistency to be factually or psychologically impossible. Or the subjectivist may need to employ the strategy of building a logical consistency requirement into the theory from the start, thus excluding from the set of value-grounding attitudes those that lead to a breakdown of consistency.

Are there any further options available to the subjectivist at this point? Perhaps the subjectivist may insist that consistency in fact doesn't always hold! If so, then there is no problem here after all. Given that the realm of value is not an objective realm—a realm of fact—but a product of our attitudes, the expectation that it be perfectly consistent may be misguided to begin with. If our attitudes are not perfectly consistent and value is a product of our attitudes, then it should come as no surprise that values may sometimes not be perfectly consistent. At least it is not something that any plausible theory should be expected to deliver. A subjectivist may insist that

¹³⁹ Dancy (2004) (p. 87)

this is as it should be, and perhaps find further support for this conclusion in the fact that our everyday evaluative thought and practice are surely not always perfectly consistent. I suspect that this is a point where different subjectivists might want to go in different directions, some insisting that consistency is an essential feature of value and a desideratum for a plausible theory of value, whilst others might say that this is an intuition that is essentially objectivist in nature and therefore has no obvious place in a subjectivist framework.

A final point: an objector might concede that subjectivism incorporates consistency and that all is well as long as we consider attitudes that are being held at the present moment. Recall what Blackburn said '[i]t is not possible to hold an attitude to a thing because of its possessing certain properties and, *at the same time*, not hold that attitude to another thing that is believed to have the same properties'. But imagine now, the objector goes on, that the subject has one attitude at one time, and a different attitude towards the same object, in virtue of the same properties, at another time. All I wish to stress here is that this objector is asking for a quite different sort of consistency than the one that I have so far been discussing. I will return to issues pertaining to time below.

Some Concerns Regarding Alm's Proposal

Having recognized that there is more than one route for the subjectivist to take when it comes to accounting for consistency, I believe that Alm's proposal is the most promising. But though attractive in its own right, this proposal raises a few concerns.¹⁴⁰ One concern focuses on the implication that an object that has never been the object of the subject's attitude may nonetheless be valuable relative to the subject in question. Does not this implication undermine the very idea of subjectivism, that things are valuable only by being objects of subjects' attitudes? Is it not odd for someone who holds that all values are grounded on attitudes to claim that an object that is

¹⁴⁰ As David Alm himself has pointed out to me, his proposal seems to presuppose that properties exist as universals. A subjectivist who for whatever reason does not acknowledge universals might therefore need to go some different route.

not, never has been and perhaps never will be the object of an attitude may nevertheless carry value?

One way to think about this would be to distinguish between direct and indirect value constitution. This distinction, or something like it, is present in Gauthier. He writes,

Of course value may be ascribed to an object—an apple, for example—or more properly to a state of affairs involving an object—eating an apple—without supposing that the object or the related state of affairs does in fact enter into the preferences of any person. The apple may be sufficiently similar to others actually involved in human preferences to permit an inference to the hypothetical preferences of some person concerning states of affairs in which it might be involved. And in this way value may be assigned to those states of affairs. But objects or states of affairs may be ascribed value only in so far as, directly or indirectly, they may be considered as entering into relations of preference.¹⁴¹

On Alm's proposal, like that of Gauthier, an object that has never itself entered as an object of any of the subject's attitudes may still be valuable relative to the subject. But, unlike that of Gauthier, on Alm's proposal it is not because the subject would have favoured the object (and that we know this because it is similar to an object that we know has actually entered as an object of his or her attitude) that we can appropriately assign the same value to the two objects relative to the subject in question. Instead, it is because the properties of objects that actually enter the subject's attitudes are thereby conferred the status of being value-makers relative to the subject. Hence, an object that has never itself entered as an object of any attitude of the subject may nevertheless be valuable relative to the subject in virtue of having properties that have been conferred the status of being value-makers relative to the subject. The value of an object may thus be said to be directly grounded on an attitude when the object is the object of the subject's attitude, and the value is indirectly grounded on an attitude when the object

¹⁴¹ Gauthier (1986) (pp. 46–47). My emphases.

is valuable in virtue of possessing properties for which the subject actually favours some other object.¹⁴²

One might object that drawing the distinction between direct and indirect value-constitution amounts to no more than saying the same thing in slightly different words and that it thus fails to address the initial concern. But it is important to note that even though something that is not an object of an attitude may nonetheless be of value, Alm's proposal does not amount to abandoning the idea that value is grounded on attitudes. Far from being at odds with the core idea that value is grounded on attitudes, the consequence that things that are not themselves objects of attitudes may nevertheless have value is a natural implication of the core idea given Alm's proposal. Since the status of being value-makers is always conferred by actual attitudes, all values are still held to be constituted by actual attitudes, even if some of these values accrue to objects that have not themselves been the objects of attitudes.

Perhaps there are some subjectivists who would want to deny indirect value-constitution and insist that only objects that have actually been objects of attitudes have value. Whilst this cannot be excluded, I don't think that there are any good arguments for such a restrictive view. I believe that it is only natural for subjectivists to recognize indirect value-constitution, and it should be seen as a virtue of Alm's proposal that it accounts for indirect value-constitution in such a natural way.

Another potential problem with Alm's proposal is that it might lead to an objectionable over-intellectualization of valuing in requiring subjects' attitudes to be very complex. It is not always and may in fact seldom be the case that we are aware of the whole set of properties in virtue of which we favour or disfavour something. But on the view we are considering, the value-grounding attitude must be properly discerning in order to pick out the relevant properties of the valuable object (the supervenience base). This concern is perhaps especially urgent if, like Dancy, one distinguishes between the supervenience base of value and the features that are value-makers. The latter are only a part of the supervenience base, which also contains various contextual features: enablers, disablers, etc. To do justice to

¹⁴² This will only be so when the objects are relevantly similar and as noted there might be more to relevance here than only the value-making properties of the object. Enabling and disabling conditions might be thought to make a difference.

this, the subjectivist following Alm's proposal would have to say that the subject's attitude confers the status of value-maker and the status of enabler, disabler, etc., to very complex combinations of properties—so complicated, perhaps, that it is not clear that they can be included as the intentional content of an attitude.

I am not sure how to adjudicate this concern, especially since I am not aware of any reliable way to measure the degree of complexity of an attitude. One potential subjectivist response to the charge of over-intellectualization would, however, start by noting that even though we may not always and perhaps may seldom be aware of or able to articulate precisely what makes something valuable, this does not necessarily mean that our attitude is not sufficiently discerning. The attitude may well be sensitive to both that in virtue of which an object is favoured and that which enables (or disables) these features to be good-making relative to us. This may be so even when we are not able to consciously pick out or articulate to ourselves or to others precisely the features that make the object good or bad, or what enables these features to play this value-making role. One possible way to highlight the relevant features and bring them into sharper focus would be to imagine the valuable object in various alternative contexts or in situations in which it gains or loses some of its features. This would be an intellectual activity for sure, and one in which we all sometimes engage, but it is not one that subjectivists require valuing subjects to engage in all the time. We should also note that it is not required of the valuing subject to possess the concepts of value-maker, enabler, etc., or even that of goodness and badness itself, in order to ground a value. What is needed is for the subject to have a sufficiently discerning attitude.

VI. Is Value-Constitution a Conscious Act?

As argued above, when we imagine hypothetical cases in which we have different attitudes than those that we in fact have, it is our actual attitudes regarding these hypothetical cases that ground values. A parallel case could be made concerning actual situations in the past or the future. When we consider situations in our actual past in which we had different attitudes from those that we have at present, it is our present attitudes that ground

values. Consider an attitude that you had in the past but that you no longer have. Let's say you used to want X for its own sake, whilst now you do not want X at all. Is the subjectivist committed to saying that X was valuable in the past but no longer is so? Or, when considering an attitude that you will come to have in the future, that X will be valuable in the future but is not yet so in the present? No. If your present attitude towards X extends to the past or the future and is not sensitive to your past or future attitude, then it is not the case that X was valuable in the past relative to you now, nor that it will be valuable in the future relative to you now. It is your present attitudes for the past or future that determine the values of past or future objects, not the attitudes that you used to but no longer have, or the attitudes you will come to have. Like merely possible attitudes, actual past and future attitudes are relevant only insofar as your present attitudes are sensitive to them. In the type of cases described, the subjectivist is committed only to non-evaluative claims like 'X is good-relative-to-you-in-the-past' or 'X is good-relative-to-you-in-the-future', not to the evaluative claim that X was or will be good.¹⁴³

A potential concern here is that we don't really have all that many attitudes at one and the same time, which would result in an axiology that is very poor in content. This would probably be so if only the attitudes of which we are presently conscious are constitutive of value, but not if we have a multitude of attitudes at any one moment, most of them unconscious. I turn now to the issue of whether an object is presently being the object of conscious attention makes a difference to value constitution, and if so, how.

Rønnow-Rasmussen has suggested that a subjectivist needs to choose between two main options with respect to value-constitution and time. According to what he calls the coexistent view, 'the value and the constitutive act cannot in time exist apart from each other.'¹⁴⁴ This description of the coexistent view is accurate as far as it goes, but one potentially misleading aspect of it must be pointed out. I have insisted that an attitude and the object of that attitude can be far separated in time. For

¹⁴³ In a different sort of case it might well be that something was or will be good relative to you now. For example, Genesis used to be a great band in the 1970s, and being able to cure cancer will be a good thing. In these cases it is not a change in attitudes over time, but a change in the objects of value.

¹⁴⁴ Rønnow-Rasmussen (2003) (p. 255)

example, I can here and now wish for its own sake that there will be human colonies on other planets two hundred years from now. In that case the value and the constitutive act might seem to exist apart from each other in time. But this is so only in the sense that the object upon which the value supervenes is located in another time. The value and the attitude on which the value is grounded are still coexistent in the relevant sense: that there will be human colonies on other planets in two hundred years is good relative to me now and will continue to be so as long as I retain my attitude.

The alternative to the coexistent view is claimed to be the 'valuable until further notice' view, according to which 'once a constitutive act has taken place there is value until the subject changes his mind'.¹⁴⁵ But how, more precisely, are we to understand these two views? Rønnow-Rasmussen asks, 'Will the valuable object cease being valuable at the literal instant the subject turns his attention to something else?' Is an object valuable only so long as it is the intentional object of the subject's conscious attitude? The proponent of the coexistent view is thought to answer yes, but if the answer is no we are presumed to end up with the 'valuable until further notice' view. However, as phrased the question is misleading. It seems to presuppose that in order for an object to be valuable it must either continuously be, or at some point have been, an object of conscious attention. But there is no need for the subjectivist to accept this presupposition. A constitutive act could very well be thought to be going on in the background, without the valuable object entering, or even ever having entered, into the conscious attention of the subject. Someone could, for example, be said to have a positive final attitude towards his or her own continued existence even if this has never been the object of her conscious reflection. It *might* be brought to the conscious attention of the subject, but it need not be so, or ever have been so. (Perhaps the subject in question leads a particularly safe and unreflective life.) This would not make her survival any less valuable relative to her.

Let us first consider the 'valuable until further notice' view more carefully. On this view, the value of an object is constituted at the moment the subject forms the relevant attitude towards that object and persists until the subject 'changes his mind'. If it is assumed that both forming an attitude and changing one's attitude are conscious acts, then this view seems to have

¹⁴⁵ Ibid (pp. 255–256)

the following consequence. If a subject, S, forms an attitude towards an object, O, at time t_1 , and O never again enters S's conscious reflection, the value of O persists indefinitely. This would seem to be so even after S has died or entered into a permanent coma, for example. Perhaps we might welcome the consequence that the values constituted by a subject can persist after the subject's death, but what would it even mean for values to persist in this manner? This seems to amount to saying that there is now a value that was constituted by an attitude that nobody has anymore. But given the idea that values can only be seen from the first-person perspective of the subject who has the relevant attitude, this would mean that there is now no perspective from which this value is accessible.

Interpreted in this way, the 'valuable until further notice' view seems to invite a literal interpretation of value projection, according to which our attitudes literally add something to the world. Such a view will be discussed further in Chapter 7 and found wanting. If we adopt the 'valuable until further notice' view, we must also explain how a value might cease to persist. Can literally projected values somehow be retracted? If so, would this mean that whatever was added to the world somehow gets 'sucked back in', or instead somehow annihilated when the subject forms a new, neutral attitude towards the same object?

A further consequence of the 'valuable until further notice' view, given that it presupposes that value constitution must be a conscious act, is this: if an object enters into S's conscious attention at time t_1 , and an attitude A_1 is formed towards this object, and then the same object re-enters S's consciousness at t_2 , when a different attitude, A_2 , is taken towards it, then the value constituted at t_1 must be seen as having persisted between t_1 and t_2 . This would be so no matter how long the interval between t_1 and t_2 —even if the subject was 12 years old at t_1 and 82 years old at t_2 . This might not seem problematic in its own right, but what if the attitude had quietly withered away in the background already before age 22 without entering into the conscious attention of the subject? Must the proponent of the 'valuable until further notice' view then say that the value nonetheless persists until S consciously contemplates the object the second time at age 82? If so, we

once again seem to end up with a value that is grounded on an attitude that nobody has and that is accessible from no perspective.¹⁴⁶

If we drop the assumption that the valuable object must be, or must at some point have been, an object of conscious attention, however, the ‘valuable until further notice’ view seems to collapse into a form of the coexistent view: the value of O_1 persists as long as S_1 ’s attitude towards O_1 persists, whether or not O_1 is, or ever has been, the object of S_1 ’s conscious attention. That is, the value is there as long as, and only as long as, constitution is going on, which it is as long as the subject in question possesses the relevant attitude, whether or not the subject is consciously aware of having that attitude. This is the interpretation that I will follow.¹⁴⁷

¹⁴⁶ One could also consider in this context a different subject S_2 who—in accordance with the participant model that will be discussed in Chapter 5—identifies with S_1 when he is 42 years old and thus takes over S_1 ’s attitudes at that time. A_1 would then not be among the attitudes taken over from S_1 and neither would A_2 .

¹⁴⁷ Rønnow-Rasmussen has pointed out to me in personal communication that, even though this is not made explicit in his (2003) essay, the discussion there was predicated on a distinction between occurrent and dormant attitudes and on the assumption that it is only the former attitudes that can constitute value. My arguments in this subsection are intended to cast doubt on this assumption. Ingemar Persson (2005) makes the distinction between occurrent and dormant attitudes with reference to conscious attention: ‘[O]ur (occurrent) attitudes and actions are determined by what we episodically think of, or attend to . . . But we are also said to possess desires, although our minds are occupied with other matters. I can be said to have a desire to continue working on this manuscript at times when all thoughts of it are far from my mind, and indeed even when I am asleep or unconscious. Such desires are *dormant*’ (p. 55). I think that subjectivists—insofar as they acknowledge this distinction and understand it in anything like this way—should say that dormant attitudes also constitute value. In Persson’s example, the value that accrues to continuing to work on his manuscript is grounded on a dormant desire.

4. WHICH ATTITUDES? MOTIVATION AND IDEALIZATION

It is characteristic of living mind to be *for* some things and *against* others. This polarity is not reducible to that between 'yes' and 'no' in the logical or in the purely cognitive sense . . . To be 'for' or 'against' is to view with favor or disfavor; it is a bias of the subject toward or away from. It implies . . . a tendency to create or conserve, or an opposite tendency to prevent or destroy. This duality appears in many forms, such as liking and disliking, desire and aversion, will and refusal, or seeking and avoiding.

Ralph Barton Perry, *General Theory of Value*

I. Introduction

The defining claim made by value subjectivists—and denied by objectivists—is that value is grounded on attitudes. But precisely which attitudes are taken by subjectivists to ground values? I have been assuming that the relevant attitudes are attitudes of favouring and disfavouring, where these terms are placeholders for a wide range of different pro and con attitudes, including desire, preference, love, respect, admiration, and

more.¹⁴⁸ Notably absent from this picture are cognitive attitudes such as beliefs. But we might ask on what basis subjectivists deny cognitive attitudes the same value-grounding powers that they grant to pro and con attitudes? Shakespeare's Hamlet declared that 'there is nothing either good or bad, but thinking makes it so'. Commenting on Hamlet's dictum, David Gauthier points out that 'thought is not the activity that summons value into being' and that 'desire, not thought, and volition, not cognition, are the springs of good and evil'.¹⁴⁹ But the question remains why this is thought to be so.

The first main objective of the present chapter is to offer a modest rationale for this restriction to pro and con attitudes. This rationale centres on the conative nature of these attitudes. The attitudes that are taken to ground value are restricted to those that involve a motivational component: some impulse, striving or effort. Purely cognitive attitudes such as beliefs are often assumed to lack any such conative aspect. '[R]eason alone can never be a motive to any action of the will', as David Hume put it.¹⁵⁰ Some philosophers have denied this, maintaining that beliefs alone can be sufficient for motivation without the help of any desire or other non-cognitive attitude.¹⁵¹ Nothing I say in what follows is intended to convince these philosophers otherwise.

Since subjectivism is here treated as a theory of the nature of value, it is not necessarily committed to any particular theory of the nature of motivation. However, on the proposal that the value-grounding attitudes should be identified with those attitudes that have a motivational aspect, value subjectivism would seem to entail some form or forms of what is known as motivational internalism. I attempt to specify which forms of internalism can or must be held in conjunction with value subjectivism. My purpose in discussing internalisms is to illustrate the implications of the

¹⁴⁸ It follows from how I have interpreted subjectivism (the object interpretation) that only affective states that have objects count as value-grounding. Any type of affective state that lacks intentional objects cannot be grounding value.

¹⁴⁹ (1986) (p. 21)

¹⁵⁰ *Treatise* (ii.iii.iii.)

¹⁵¹ One example is Scanlon (1998) (pp. 37–41). Some hold that neither beliefs nor desires are necessary for motivation. On Jonathan Dancy's (2000) view, it is facts that motivate, not beliefs about facts.

subjectivist theory: what value subjectivists are and are not committed to in this regard.

The chapter's second main objective is to consider ways in which a subjectivist might want to further restrict the class of value-grounding attitudes or otherwise idealize these attitudes. I have previously argued that it is only the attitudes that we actually have, as opposed to the attitudes we could have had but don't actually have, that are taken to ground value. The attitudes of hypothetical subjects, I have maintained, are relevant only insofar as the attitudes of actual subjects are sensitive to them (which they sometimes are and sometimes are not). It should be remembered that the objects of our actual attitudes need not themselves be actual. The constitutive ground of a value is always identified with some actual attitude, but the bearer of the value (i.e. the object of the attitude in question) might be located in a different possible world. It is the attitudes of actual subjects that constitute the value of both actual and hypothetical objects. Recall that whilst value-constitution is an expression of a particular world-bound perspective—the perspective determined by the attitudes we actually have—values are not world-bound in their range of application.

Value-grounding attitudes are further restricted to those that are held at the present time, and I have argued that the object of an attitude retains its value relative to the relevant perspective for as long as and only as long as the attitude in question is retained, even though the valuable object need not have been an object of conscious attention. However, the restrictions to actual and present attitudes are not sufficient to identify the relevant value-grounding perspective. In the next chapter I will distinguish two different versions of value subjectivism that I label individualist relativism and universalist absolutism. On the former, the relevant value-constituting perspective is that of the individual subject as determined by that individual's present and actual conative attitudes. The attitudes of other actual subjects enter the picture insofar as the attitudes of the individual in question are sensitive to them. On the individualist relativist view, then, there is more than one value-constituting perspective in one and the same possible world at the same time, with each such perspective giving rise to its own distinct set of values that may or may not overlap with those of others.

When we exclude attitudes that are not actual, not present, not conative, and—like individualist relativists—not the individual subject's own, this does not amount to idealization. Universalist absolutism, on the other hand, does involve a strong form of idealization. On this view, the relevant value-grounding perspective is that of an idealized 'participant': an

actual, individual subject who has identified with all other actual subjects and made their attitudes her own. No values are thought to be constituted by the individual subject until after the subject has gone through this process of universal attitude takeover. Through this universalization mechanism, diverse and conflicting attitudes of separate individuals are transformed into a single, unique set of values relative to the actual world. This view still deserves to be called a version of subjectivism, as it holds that values are grounded on and constituted by attitudes. But we might ask whether the rationale for this form of idealization is not in tension with the basic subjectivist idea. I argue in the next chapter that there is a conditional presumption in favour of individualist relativism. That is, insofar as we are subjectivists, we ought to go for the individualist version of subjectivism.

In the present chapter I concentrate on more modest forms of idealization that can at least potentially be endorsed by subjectivists of both individualist relativist and universalist absolutist persuasion. The question to be examined in this context is whether sets of value-constituting attitudes should be restricted, or otherwise idealized in some way or other, and, if so, on what basis. Some philosophers have argued that idealization as such, even of modest kinds, is incompatible, or at least in tension, with the very idea behind subjectivism. As David Sobel characterizes them, these critics claim that ‘the only good rationales for such idealization are incompatible with the animating spirit of subjectivism. The only legitimate rationales for idealizing the attitudes that are granted authority . . . are rationales that presuppose a nonsubjectivist grounding of values.’¹⁵² Following Sobel, I offer a way in which a subjectivist could idealize whilst remaining true to the spirit of subjectivism.

II. Values Are Not Grounded on Beliefs

Why should we think that it is those (and only those) attitudes that involve being for or against something that ground values? Insofar as one allows that

¹⁵² Sobel (2009) (p. 338)

beliefs are motivationally inert, as David Hume believed, it is plausible to think that it is not beliefs that have the power to ground values, but that this power is the exclusive domain of 'active' attitudes such as desires. But this is only the beginning of an answer. A slightly different way to approach the issue is to focus on the 'polarity' of pro and con attitudes and claim that this makes them uniquely suited to ground values. After all, positive attitudes (attitudes of favouring) are taken to ground positive values, and negative attitudes (attitudes of disfavouring) to ground negative values. Moreover, when some object carries a positive value relative to some subject, that subject is (at least normally) motivated to respond to the object in some affirming way: to choose it, create it or preserve it, for example. And in the case of objects of negative value, the subject has a corresponding motivation to respond in some contrary manner: to abolish, destroy or prevent it.

Intuitively, this polarity does not easily reduce to that between truth and falsehood. Hume assigned different roles to reason and to taste respectively. Reason, on Hume's view, is the faculty that 'conveys the knowledge of truth and falsehood' and 'discovers objects as they really stand in nature, without addition or diminution'. Taste, by contrast, is what 'gives the sentiment of beauty and deformity, vice and virtue' and 'has a productive faculty, and gilding or staining all natural objects with the colours, borrowed from internal sentiment, raises in a manner a new creation', that is, creates value.¹⁵³

One might perhaps insist that the positive/negative polarity could be accounted for in terms of beliefs that an object deserves to be favoured or disfavoured. But this proposal is unsatisfactory for several reasons. The polarity that is captured by such beliefs about merited attitudinal responses is entirely parasitic on the polarity of the pro and con attitudes that are believed to be merited. This may not be a problem for the proposal considered in its own right, but one wonders why we should want to bring in beliefs in this roundabout way. Does bringing in beliefs in this way accomplish any work in the theory? If not, the appeal to beliefs seems unnecessary. But if we think it might fill a theoretical function, then what could that function be? If it is to bring in objectivism through the backdoor

¹⁵³ *Enquiry* (appendix I)

by appealing to an attitude-independent standard of merited or fitting attitudes, then it clearly begs the question against the subjectivist.

An obvious difference between cognitive attitudes such as beliefs and conative attitudes of the pro and contra type is that the former can be true or false whilst the latter cannot. Moreover, we can all agree that there is an objective standard of truth and falsity of a belief whilst whether there is any corresponding standard with respect to the fittingness of pro and con attitudes is a matter of controversy between subjectivists and objectivists. The proposal to ground values on beliefs about merited responses, whatever intrinsic merits we might think it has, does not seem to go well together with subjectivism.

III. Motivational Internalisms

The thesis that value is grounded on attitudes trivially entails that there is a necessary connection between value and attitudes. Put differently, whenever there is a value there necessarily is a constitutive ground of that value, which is always some attitude. If now the attitudes that make up constitutive grounds of value are taken to be precisely the attitudes that involve a motivational component, as I propose here, it follows that there is a necessary connection between value and motivation. Coupled with this proposal, value subjectivism thus entails some form of what is known as motivational internalism. But what forms exactly?

Judgment Internalism

Consider first what is commonly known as judgment internalism. On this form of internalism, there is a necessary connection between thinking or sincerely saying that something is good, on the one hand, and motivation on the other. This necessity may be conceptual or metaphysical, but the underlying thought is that it is not merely a contingent matter that people are motivated to pursue what they think or sincerely say is good.

Should some form of judgment internalism be attributed to the value subjectivist? The answer may depend on which theory of evaluative language

that is adopted in conjunction with a subjectivist theory of value. As noted in Chapter 1, value subjectivism, as I understand it, is neutral between cognitivist/descriptivist and non-cognitivist/expressivist theories of evaluative language. On a view of the latter type, S's judgment that O is good is thought to express S's favouring of O, and this is commonly assumed to imply that, necessarily, S is motivated. This bridge from non-cognitivism/expressivism to motivational judgment internalism has been questioned, however. The issue depends on how precisely we are to understand the expression relation.¹⁵⁴

On cognitivist views, S's judgment that O is good is a belief that O is good and as such it does not necessarily imply that S is motivated. A cognitivist who is also a subjectivist could think that when a given subject believes something to be valuable, there is always a motivating attitude in the background playing a causal role in producing the subject's belief. And if beliefs about value, whether true or false (even if systematically false, as on the error theory) are always caused by attitudes in this way, then this seems to imply a necessary (though causal rather than conceptual) connection between value judgment and motivation. Whether such a cognitivist subjectivist can be a judgment internalist is a question that I will leave open. Since value subjectivism by itself does not imply judgment internalism—and it remains unclear whether value subjectivism in conjunction with either cognitivism or non-cognitivism does—I can leave this form of internalism aside.

Knowledge Internalism

The view that John Mackie ascribes to the Plato of the Republic is another form of motivational internalism. On this view, which might be called knowledge internalism¹⁵⁵, there is a necessary connection between knowing what is good and being motivated to pursue it. Goodness motivates whoever happens to be acquainted with it. Mackie writes,

¹⁵⁴ See Joyce (2002) (pp. 337–339)

¹⁵⁵ This label was suggested to me by Wlodek Rabinowicz.

In Plato's theory the Forms, and in particular the Form of the Good, are eternal, extra-mental, realities. They are a very central structural element in the fabric of the world. But it is held also that just knowing them or 'seeing' them will not merely tell men what to do but will ensure that they do it, overruling any contrary inclinations. The philosopher-kings in the Republic can, Plato thinks, be trusted with unchecked power because their education will have given them knowledge of the Forms. Being acquainted with the Forms of the Good and Justice and Beauty and the rest they will, by this knowledge alone, without any further motivation, be impelled to pursue and promote these ideals.¹⁵⁶

Or, as Stephen Darwall describes Plato's view, 'it is of the nature of the Good that it cannot be fully grasped with indifference'.¹⁵⁷ Since one cannot have knowledge without judgment; since one cannot know that something is good without believing it to be so, knowledge internalism is close to judgment internalism. But the necessary connection postulated by knowledge internalism is not between motivation and any belief about value, but between motivation and specifically beliefs that are true and justified. So, though it is not the case that value is taken to motivate us by its sheer existence alone (since it must be grasped), knowledge internalism does attribute intrinsic motivational powers to goodness and badness as such.

One major problem with Plato's view as described is that it cannot explain why this is so: how can goodness and badness, seen as 'eternal, extra-mental, realities' necessarily motivate anyone who happens to be aware of them? Mackie argues that value properties with intrinsic motivational powers are too 'queer' to be countenanced in a plausible ontology. Darwall concurs: '[I]t is mysterious how there could be such a property, precisely because it seems to require a bridge between knowledge, or its object, and action.'¹⁵⁸ A further problem is that, since knowledge requires objectivity, knowing that something is good seems to require that it is *objectively* good.¹⁵⁹ And since I am now investigating what kinds of internalism that are

¹⁵⁶ Mackie (1977) (pp. 23–24)

¹⁵⁷ Darwall (1992) (p. 157)

¹⁵⁸ Ibid (p. 158)

¹⁵⁹ In Chapter 8 I briefly discuss truth and knowledge.

implied by or fit together with subjectivism, knowledge internalism is not what I am looking for.

Nevertheless, the subjectivist finds something essentially right in Plato's internalism. It is not merely that it captures something true about the phenomenology of (what is often called) 'moral motivation': the fact that it may appear to us as evaluating subjects as if being acquainted with goodness impels us to pursue it. It is also that subjectivism implies that, in a certain kind of way, this is actually so! Indeed, subjectivists could be argued that they can hold on to what is appealing in Plato's view, whilst at the same time be able to explain motivation in a way that is not available to an objectivist like Plato. If values are grounded on specifically conative attitudes it should not be surprising that they possess motivating powers. If values are projections of attitudes that invariably involve a motivational component, then it naturally follows that there is a necessary connection between that which is valuable relative to subjects and their motivation to pursue it. However—and herein lays its superior explanatory power—on subjectivism it is not strictly speaking the goodness or badness itself that is motivating, but rather the attitude on which the goodness or badness is grounded.

It should be remembered that the motivating attitudes in question can figure entirely in the background and need not appear at all in the content of the subject's deliberation. When we ask 'What motivated you to pursue this good?' we should not in general expect the answer to mention an attitude, but rather a list of the features of the chosen action for which it was chosen. Recall Stephen Darwall's observations that 'the locus of the desiring person's concern is most typically the object [of her desire] and the facts connected with it rather than the fact that she has the desire' and that '[f]rom the fact that any intentional action must be explainable by the agent's desires it does not follow that the agent's reasons for acting must itself make reference to his desires'.¹⁶⁰ From the valuing subject's first-person point of view, what matters are the objects of her conative attitudes and the features of these things for which she favour or disfavour them. But when we want to explain the subject's actions from the outside, from the third-person point of view, her attitudes must be mentioned. Simon Blackburn has argued that it is a virtue of a subjectivist (projectivist) view that it is

¹⁶⁰ Darwall (1983) (p. 37)

readily compatible with the ‘standard model of explanation of why someone does something’. This model attributes both a belief and a desire to the agent: ‘The belief that a bottle contains poison does not by itself explain why someone avoids it; the belief coupled with the normal desire to avoid harm does’.¹⁶¹ Expanding on Blackburn’s example, if harm is the object of a subject’s negative attitude, harm is bad relative to the subject in question. If we ask, ‘What motivated you to avoid this harm?’ the subject’s answer might well be ‘that harm is bad’. But that harm is bad need not be appealed to in the explanation of the action, the explanation need only involve the subject’s beliefs that the bottle contained poison and that poison is harmful, plus her negative attitude towards harm. Even if not necessarily committed to it, the value subjectivist can readily accept this highly plausible theory of motivation.

Existence Internalism

A third type of motivational internalism is what Stephen Darwall calls existence internalism.¹⁶² This view concerns not what it is to accept or express a value judgment, but what it is for such a judgment to be true. As with knowledge internalism, the thought here is that it is not merely a contingent matter that people are motivated to pursue what is good. But whilst on knowledge internalism value ‘is such that motivation is necessarily an effect of engaging it epistemically’, on existence internalism a necessary relation holds between knowledge of value (‘at least when it is gained by an agent for herself within the deliberative process of practical reasoning and judgment’) and motivation because ‘practical knowledge is *of motives*’.¹⁶³

¹⁶¹ The argument is spelled out in Blackburn (1984) (pp. 187–189).

¹⁶² Darwall (1992). Darwall uses Plato’s view to exemplify a version of existence internalism on which value ‘is such that motivation is necessarily an effect of engaging it epistemically’. This form of existence internalism is contrasted with a very different form of existence internalism according to which motivation is taken to be ‘a constituent of ethical facts themselves’ (p. 157). Darwall calls the latter ‘constitutive’ and the former ‘non-constitutive’ forms of existence internalism. Note that in doing so Darwall is using the term *constitution* in a very different sense from the sense in which I am using it.

¹⁶³ Ibid (p. 158). Emphasis in original.

The underlying objective of existence internalism is to understand normative force in terms of motivating force. This project forms no part of subjectivism.¹⁶⁴ Motivation is straightforwardly an objective phenomenon; there are objective facts about people's motivations. Hence, on existence internalism it is presumably possible to determine, from a third-person point of view, what is good or bad by looking at facts about motivations. On David Sobel's characterization of 'subjective accounts of well-being', for example, 'one's rationally contingent nontruth assessable proattitudes ground true claims about what is good for one'.¹⁶⁵ This is not possible on subjectivist views as I understand them here. On such views, value, or 'normative force', is not something objective. Goodness and badness are products of our attitudes and can only be seen from within the various first-person perspectives of individual subjects who have the attitudes in question, through the lens of their attitudes. This cannot ground objectively true claims about what is good for someone (or what is good *period*), but only of what is good (either for or period) *relative to* a subject's perspective, and such relative-to-claims, when true, are not *evaluative* truths.

Consider the view known as reasons internalism as one form of existence internalism. On this view, as famously defended by Bernard Williams, in order for a person P to have a normative practical reason to do A, there must be something in P's 'subjective motivational set' that would be somehow furthered by doing A, or P must be able to become motivated to do A.¹⁶⁶ In opposition to reasons internalism, it could be said that there is no necessary relation between there being normative practical reasons, or a person's having such reasons, and motivation. A person's having normative reasons to do something is independent of that person's actual or hypothetical motivations.

¹⁶⁴ In the next subsection I will introduce what I call perspective internalism. This view recognizes an entailment from 'O is valuable relative to S' to 'S is motivated to take action in relation to O'. Some would like to treat this as a form of existence internalism, but given that claims of the form 'O is good relative to S' are not seen as evaluative/normative claims, I think that perspective internalism should be seen as making an internalist claim that is distinct from that made by existence internalism which posits an entailment relation between evaluative/normative claims and claims about motivation.

¹⁶⁵ Sobel (2009) (p. 336)

¹⁶⁶ Williams (1980)

The subjectivist can, and most probably will, agree with this denial of reasons internalism. Motivations (motivating reasons) can be seen from the third-person point of view, but normative reasons cannot. Seeing normative reasons requires the first-person perspective. Hence, unlike the reasons internalist, the subjectivist does not claim that the fact that there is nothing in a person's subjective motivational set implies that there is no reason for her. Indeed, no objective normative conclusions of either a positive or negative nature can be drawn from facts about the presence or absence of motivation. Rather, on the subjectivist view, relative to a particular subject's first-person perspective, there can be both internal and external reasons. For example, from S's perspective, Charlie might have a reason to do A and this regardless of whether there is anything in *Charlie's* 'motivational set' that would be promoted by her doing A. If there is nothing in Charlie's motivational set that would be promoted by her doing A, this would be an external reason for Charlie. A distinction must be drawn between there being a reason for and there being reasons relative to. In the case just described, there is a reason *for* Charlie (Charlie *has* a reason) to do A, but she has this reason only *relative to S*. Relative to her own perspective, Charlie might have no such reason. The normative status of being a reason is always relative to some individual subject's perspective.¹⁶⁷ From the third-person point of view there are no normative reasons (though there are reasons of other sorts, including motivating reasons); from this detached, objective viewpoint there are only reasons relative to this or that perspective, and relative-to-claims I have insisted are not normative.

Perspective Internalism

None of the three internalisms discussed above are what the value subjectivist is committed to. A subjectivist may accept judgment internalism, but need not do so; the debate over this kind of internalism is independent of the debate between subjectivism and objectivism. Even though we can incorporate much of what is appealing in knowledge

¹⁶⁷ Note that I am not (just) making the claim that reasons are always reasons *for* someone or other (a claim that objectivists and subjectivists alike can endorse), but the claim that a reason (for someone) is always relative to some perspective.

internalism, this form of internalism becomes strictly speaking irrelevant for subjectivists since knowledge requires objectivity. Existence internalism implies that all valuable objects necessarily involve motivation; facts about values always at least in part supervene on facts about motivations.¹⁶⁸ That value always supervenes on motivating attitudes is not what subjectivists are arguing. As I have argued, subjectivists hold that value may supervene on any object that can be the object of a motivating attitude, regardless of whether the object in question itself involves motivation.

As hinted at above, there are both important differences and some similarities between these three forms of internalism (judgment, knowledge and existence internalism) and the internalism of the value subjectivist who identifies the value-grounding attitudes with those that involve a motivational component. The latter kind of internalism, which I call perspective internalism, can be traced in what Darwall ascribes to Thomas Hobbes:

All deliberation begins in an agent's desires, but this does not mean that they begin in a premise about her desires. The deliberating agent reasons from a premise she accepts in having a desire, not from the premise that she has a desire. And this premise is something normative—that something would be good, that she is to or ought to do something. The agent has these normative thoughts because she has desires. They are the 'appearances' of her desires. As it happens, there is nothing in the nature of the objects of her desires that answers to the normative properties she attributes to them in having desires. But Hobbes evidently believes that theoretical knowledge of this fact need not undermine deliberation. So long as the agent has appetites and desires, like the desire for self-preservation, that are entirely independent of theory, and beliefs about what will accomplish their objects, she will take action in the thought that in so acting she is achieving good and doing as she ought.¹⁶⁹

¹⁶⁸ This view bears some similarities with the satisfaction interpretation of value preferentialism that I have argued is best seen as a first order view with objectivistic second order implications.

¹⁶⁹ Darwall (2000) (p. 333)

There are several things to note here. First, Darwall's Hobbes here relies on the background/foreground distinction that I introduced in Chapter 2 according to which a desire, or some other pro or con attitude, can have a motivational presence in the subject's decision-making without appearing in the content of her deliberation. Second, on Darwall's interpretation, Hobbes implicitly works with a distinction between first- and third-person perspectives, or as Darwall describes them, 'the perspective of an agent deliberating about what to do' and 'an observer's standpoint', respectively. That I actually desire something, that I have this motivating attitude, is something that can be seen from the outside, from the third-person perspective: 'Facts about my desires seem to be psychological descriptions of my situation as they might be viewed from an observer's standpoint.' It is such facts that existence internalism seeks to connect to value and normativity. On Darwall's interpretation of Hobbes, by contrast, that something is good, that it is something people have reason to desire or promote, is something that can only be seen from the first-person perspective: it concerns what there are normative reasons to do 'as these might be viewed from the perspective of an agent deliberating about what to do'.¹⁷⁰

The form of internalism we have arrived at posits a necessary connection, not between what is good and motivation, but between what is good-relative-to-a-perspective and motivation. Specifically, between what is good-relative-to-the-perspective-of-a-subject-S and S's motivating attitudes. I use the expressions 'relative to' and 'perspective' in somewhat special senses. What 'good relative to' means will be discussed in more detail later on. For the moment it will be enough to note that for something to be good relative to S's perspective is for it to be an object of one of S's pro attitudes. 'Good relative to' should thus not be confused with 'good according to'. To be good relative to is not the same as being good from what we might call the epistemic or cognitive perspective of the subject, what she believes or judges. But rather her perspective as determined by her pro and con attitudes.

Consider now the following formulation of perspective internalism:

¹⁷⁰ Ibid (p. 316)

Whenever an object *O* has some value relative to the perspective of a subject *S*, then *S* is necessarily motivated to take action in relation to *O*, at least to some degree.

This will be modified somewhat later on in the light of my discussion of idealization, but this formulation will do for now. Which action a subject will be motivated to take in a particular case depends on the value and the type of object (value bearer) involved.¹⁷¹

It should be noted that the view I call perspective internalism is in one respect weaker than the internalist view that Darwall identifies in Hobbes. In the passage quoted above it is claimed that something normative, that something would be good or that it ought to be done, must enter into the content of the subject's deliberation. On the formula just offered, if an object *O* is valuable relative to a subject *S*, then *S* is motivated to take action in relation to *O* regardless of whether *S* judges or believes that *O* has this value. Neither facts about motivating attitudes nor believed facts about value need enter into the content of the subject's deliberation. The content of the subject's deliberation can consist wholly of premises about the non-normative/non-evaluative characteristics of the options under deliberation. That is, these premises can make reference exclusively to the non-evaluative features of the options for which they are being favoured or disfavoured by the subject: that is, the properties that make up the supervenience bases of the values in question. In other words, it is not necessary that the subject judges the option desirable, it is enough that she desires it.

It is undoubtedly often the case that the subject not only favours or disfavours an object in virtue of some of its non-evaluative properties but also at the same time judges the object to be good or bad, better or worse, desirable or undesirable, etc. This, however, does not seem to be necessary either for the attitude in question to ground a value relative to the subject or for the subject to be motivated to take action in response to the favoured

¹⁷¹ Examples of actions that a subject can be motivated to take include choosing, creating, promoting, procuring, prolonging, protecting, preserving, etc. (for objects of positive value), and abolishing, avoiding, destroying, preventing, etc. (for negative values). When one is motivated to protect or destroy something, that something is plausibly a concrete object like a building or an artefact (not a state of affairs). And when one is motivated to prolong or prevent something, that something is plausibly an event like an experience or a war (not a concrete object).

object. Hence, as long as they have pro and con attitudes, perspective internalism could hold even of subjects that never make value judgments and even if they lack value concepts.

Note on the ‘Direction of Fit’ Metaphor

I have proposed that the attitudes that constitute values should be identified with those attitudes that involve a motivational component. An alternative proposal would be to single out the relevant attitudes by reference to the distinction between two different ‘directions of fit’. As David Sobel and David Copp put it, ‘Beliefs “aim to track the world”, while desires “aim to impose themselves onto the world”’. Beliefs and desires are thus distinguished by reference to the different functional roles that they are taken to play.¹⁷² A subjectivist who would take this route would identify value-grounding attitudes as those that have the mind-to-world direction of fit: those attitudes that ‘aim to impose themselves onto the world’. This wording fits nicely with the subjectivist idea that value is something that is bestowed upon or conferred on objects. However, the usefulness of the direction of fit metaphor has been questioned, and the subjectivist need not rely on it.

Note that the proposal to rely on the direction of fit metaphor in identifying those attitudes that ground value claims nothing about motivation, and as such it is distinct from the proposal to identify the value-grounding attitudes as those that involve a motivational component. Nevertheless, it seems likely that the two proposals would be extensionally equivalent. That is, both proposals would likely identify the same set of attitudes. Having access to two independent ways of identifying the same set of attitudes would speak in favour of this classification not being arbitrary. However, if it would turn out to be the case that these independent ways to identify the relevant attitudes would yield different sets of attitudes, we would have two different versions of subjectivism depending on how we choose to identify the attitudes that are taken to ground value.

¹⁷² Sobel & Copp (2001) (p. 45). The idea came originally from Elisabeth Anscombe (2000).

IV. Idealization

Ralph Barton Perry wrote, ‘Any object, whatever it be, acquires value when any interest, whatever it be, is taken in it; just as anything whatsoever becomes a target when anyone whosoever aims at it.’¹⁷³ But what if I harbour false beliefs about that at which I am aiming? Imagine a hunter taking aim at a deer-shaped rock in the belief that it is a deer. There is a sense in which the rock is now the hunter’s target, but there is also a sense in which the hunter’s aim is misdirected. It would not be out of place to insist that in one natural sense the hunter is not ‘really’ aiming to shoot the rock.

Similarly, what if the object of an attitude is not in fact the way it is imagined to be? It is easy to imagine situations where the object of a subject’s pro attitude does in fact lack some or all of the properties for which the subject favours it. I might, for example, crave the moon in virtue of its sweet taste and soft texture. In such cases we would like to say that the object is not valuable relative to me despite being the object of one of my actual and present pro attitudes. But can subjectivists go along with this without thereby abandoning their theory? Objectivists can appeal to an objective, attitude-independent standard of good and bad, but what can subjectivists appeal to in its place?

The answer, I believe, starts from Alm’s proposal, introduced in the previous chapter, that by favouring an object in virtue of certain features that this object is believed to possess, the subject bestows upon these features the status of being value-makers relative to the subject in question. As a consequence, if the favoured object in fact lacks the features for which it is favoured, it is not valuable relative to the subject despite being the object of her actual and present pro attitude. The conclusion that naturally emerges is that an object that lacks some or all of the properties for which it is favoured does not carry the value in question. After all, if values are supposed to supervene on those properties of the object for which it is favoured, and the object doesn’t actually have these properties, then the supervenience base is absent (or is not located where the subject thinks it is). There is thus nothing there for the value to supervene on. An object cannot be made

¹⁷³ Perry (1926) (pp. 115–116)

valuable by some features that it does not possess (unless, of course, it is in virtue of the absence of these features that it is favoured).

Imagine that you have a desire to drink the contents of the glass in front of you in the belief that the glass contains gin and tonic.¹⁷⁴ Contrary to what you believe, however, the glass does not contain gin and tonic but rather petrol. It would clearly be rational of you to drink the contents of the glass. After all, you believe that what you have in front of you is a perfectly ordinary cocktail. But does your desire for a gin and tonic bestow value on the object in front of you, or on the state of affairs in which you empty this particular glass? Would drinking the contents of this glass be of value relative to you? The answer is no, and the reason is because most or all of the features for which you desire to drink a gin and tonic are not present. The features in question (the refreshing taste of the tonic, the pleasant effects of the alcohol, etc.) are still value-making relative to you, and drinking a gin and tonic is still good, but this value is not located where you think it is. Your desire is 'misdirected', though it is really your belief that is at fault and not the desire as such.¹⁷⁵

Attitudes that are thus misdirected are not robbed of their value-constituting powers. Even if the object of my attitude (call it O_1) lacks the properties for which I favour it, the attitude in question still bestows the status of value-makers on those properties (call them P_1, \dots, P_n) relative to me. A different object O_2 may in fact possess P_1, \dots, P_n , in which case O_2 (even if it has never been the object of any attitude of mine) *is* valuable relative to me. Recall the distinction that I have previously made between direct and indirect value-constitution. Any object that in fact possesses P_1, \dots, P_n is indirectly conferred value relative to me, and similarly relative to any subject who shares the relevant attitude. And this holds true not just for objects located in the actual world. Thus, these misdirected attitudes still confer value, though the values in question are not located quite where the favouring subjects think they are. These values do not supervene on the objects toward which the attitudes in question are being directed, but

¹⁷⁴ The classic gin and tonic example was first described by Bernard Williams (1980) (p. 102).

¹⁷⁵ Note that what I am arguing for here is not that only the attitudes we would retain under 'full information' ground values.

instead on those objects that do possess the relevant value-making properties.¹⁷⁶

David Sobel agrees:

I am not denying that X can be the object of my desire despite my misinformation about X. My point is that there is a mismatch between the understanding an agent has of X when she has an intrinsic desire for X that would not be sustained in light of accurate information about X and the true nature of X. In these cases I want to say that the agent's desire is not genuinely for X as it in fact is. The desire is not responsive to the true nature of its object, and in that sense, is not really for it. It is this distinction, I take it, that we are after when we say that we did not 'really' want something. It was only wanted to the extent that we were ignorant of its true nature. Our informed desires are, in a sense, more genuinely for their objects. And such desires are what we ordinarily would have referred to as what we 'really want'.¹⁷⁷

Since no actual and present pro or con attitude, even if misdirected in the sense just described, is being denied the power of conferring value, this form of idealization—if it deserves to be so called—does not restrict the set of value-grounding attitudes, but merely redirects them to those objects that really do possess the properties that have been conferred the status of value-makers. This is one rather natural way in which a subjectivist can claim that a valuing subject can be mistaken about what is good or bad relative to him or her. Note, however, that it is not the attitude as such that is claimed to be mistaken, rather the mistake stems from the subject's mistaken beliefs about the object of her attitude. As Hume put it, 'a passion must be accompany'd with some false judgment, in order to its being unreasonable; and even then

¹⁷⁶ Note again that it is not the case that the object of the subject's attitude is really just the set of properties P_1, \dots, P_n . That objects are favoured *for* (some of) their properties does not imply that it is really those properties that are favoured rather than the objects that (are thought to) possess the properties. Besides, if exclusively (sets of) properties could be objects of value-constituting attitudes, this would put up an unjustified restriction on the type of objects that could be bearers of value.

¹⁷⁷ Sobel (2009) (p. 347)

'tis not the passion, properly speaking, which is unreasonable, but the judgment.¹⁷⁸

However, this form of idealization may amount to a restriction of the set of value-grounding attitudes in some special cases. As has been argued previously, some of our attitudes may confer value on objects that are located in other possible worlds, so it need not be the case that the properties in virtue of which we favour or disfavour things are instantiated in the actual world. But what if the subject favours or disfavors an object in virtue of some property or set of properties that could not possibly be instantiated? If I love a figure in virtue of its square roundness, my attitude has not conferred a value. Or should we say that I have conferred a value upon an object located in an impossible world? Or that I have conferred an impossible value?

A perhaps more interesting issue arises when we consider cases in which the subject is not mistaken about whether a certain object possesses a particular property but is rather mistaken about the very nature of the property in question. Say that S favours O on account of P, and O really does possess P, but S has a mistaken view about what having P involves. Is this even possible? Surely S can think that she favours O on account of P, and she can say that she favours O on account of P, but can she really *favour* O on account of P while lacking an understanding of what having P involves? Has S succeeded in conferring upon P the status of value-maker relative to her in this case? I doubt it. If the subject lacks a proper understanding of the nature of a property, then it is hard to see how her attitude could be sufficiently discerning to have been able to single out this property for bestowing upon it the status value-maker; if she lacks a proper understanding of P, then it is probably not really in virtue of P that she favours O, but instead for some other property. More on this below.

Why Idealize?

Insofar as the above is thought to amount to a form of idealization, it is one that seems to follow rather naturally from Alm's proposal which I have argued for previously on independent grounds. Do we need a further

¹⁷⁸ *Treatise* (ii.iii.iii.)

rationale for this kind of idealization? If we do, there is one readily available. As Sobel points out, ‘The point of the idealization is to give an accurate understanding of what the option that one is assessing is like.’¹⁷⁹ This was understood by an objectivist such as Henry Sidgwick who wrote:

It would seem then, that if we interpret the notion ‘good’ in relation to ‘desire’, we must identify it not with the actually *desired*, but rather with the *desirable*—meaning by ‘desirable’ not necessarily ‘what ought to be desired’ but what would be desired, with strength proportioned to the degree of desirability, if it were judged attainable by voluntary action, supposing the desirer to possess a perfect forecast, emotional as well as intellectual, of the state of attainment or fruition.¹⁸⁰

Sobel follows Sidgwick in talking about ‘perfect forecasts’:

Desires that do not involve this perfect forecast are, in a sense, not actually for the option *as it is* but rather for the option as it is falsely imagined to be. Only desires formed in the light of an accurate understanding, phenomenological and otherwise, of what an option would be like are responsive to the true nature of the option under consideration.¹⁸¹

Sobel is here discussing subjectivist theories of well-being and not value subjectivism in general, but as he himself suggests, his proposal can be extended to cover other areas beyond well-being. Since I’m discussing subjectivist theories of value and not of well-being and operating not just with desires but with a wider set of pro and con attitudes, ‘what an option would be like’ needs to be broadened into something like ‘what the object is really like’, where *object* is taken to cover not just events such as experiences, but also states of affairs, concrete objects, persons, etc. ‘Perfect forecast’ needs to be similarly broadened into something like ‘accurate picture’ as not all favoured or disfavoured objects lie in the future—or, in the case of abstract objects, in any particular time. With these modifications, however, Sobel’s proposed rationale for idealization seems to suit my subjectivist’s

¹⁷⁹ Sobel (2009) (p. 345)

¹⁸⁰ Sidgwick (1981) (pp. 110–111)

¹⁸¹ Sobel (2009) (p. 345)

purpose. Sobel further insists that '[n]o independent account is needed to explain why the procedure must involve an accurate understanding of what the [objects of one's attitudes are really like]—this requirement is continuous with the subjectivist thought, not independent of that thought.'¹⁸²

As noted at the outset of this chapter, the objection has been raised that idealization as such (even of modest kinds) is incompatible, or at least in tension, with the very 'animating spirit' of subjectivism.¹⁸³ It is thus argued that the rationale for idealizing attitudes could only be to better track independent facts about what is good or bad; to track objective values. Sobel's response was that idealization is indeed an attempt to get already existing facts into better focus, but the facts in question are just the ordinary, non-normative/non-evaluative facts about the true nature of the favoured or disfavoured object (and perhaps the true nature of the properties in virtue of which the object is favoured or disfavoured), not facts about what is objectively good or bad. The modest form of idealization that I have argued for similarly does not appeal to any objective standard of value, but only to such things as were already components of the subjectivist analysis. This, I believe, should relieve us of the general worry that idealization as such is incompatible with, or in tension with, value subjectivism. This form of idealization is not only compatible with subjectivism but is continuous with it.

A Final Word on Internalism

I have argued that where an object, O_1 , is favoured by subject S for properties that it does not in fact have, and there is some other object, O_2 , that does in fact possess these properties, it is the case that O_2 is valuable relative to S , and this is so even when S is unacquainted with O_2 . This might

¹⁸² Ibid. Insofar as we find the reasoning in the last paragraph of the previous section compelling, we may want to add the following: not only do subjects need to have an accurate understanding of what the objects of their attitudes are really like; they also need an accurate understanding of what the properties in virtue of which they favour or disfavour these objects are really like.

¹⁸³ See, for example, David Enoch (2005) and Arthur Ripstein (2001).

seem to spell trouble for the internalist thesis I proposed earlier according to which whenever some object is valuable relative to S, S is necessarily motivated to respond to it. This does not seem to hold if S is unaware of the existence of the object in question. After all, how can one be motivated to pursue something one is not aware of?

In the light of this, the perspective internalist thesis needs to be reformulated along the following lines:

Whenever an object O has some value relative to the perspective of a subject S, then, insofar as S is aware of (and has an accurate picture of) O, S is necessarily motivated to take action in relation to O, at least to some degree.

V. Restrictions on Value-Grounding

Attitudes?

Are there further forms of idealization open to value subjectivists? Precisely which forms of idealization are open to subjectivists remains an open question, and this would need to be considered on a case-by-case basis. The idealization, if it deserves to be so called, that I have argued for above does not restrict the set of attitudes that are taken to ground value, except, perhaps, in special cases¹⁸⁴, it merely redirects them, when needed, to objects that really have the properties that are seen as good-making. It might be insisted that this idealization is not idealization enough, and that we must idealize further and restrict the set of value-grounding attitudes in some way. In this section I consider a few potential restrictions on value-grounding attitudes.

A Restriction to Rational Attitudes?

¹⁸⁴ As when an object is favoured in virtue of impossible properties.

David Gauthier suggests that preferences must be rational in order to give rise to values. More precisely, he imposes two rational conditions on preferences: they must be coherent and they must be considered.

If the conditions for coherent preferences are met, we may introduce a measure of the objects of preference. If the conditions for considered preference are also met, we may identify this measure with value, and its maximization with rationality. We may therefore speak of the conditions for coherent and considered preferences as conditions of rational preference.¹⁸⁵

That preferences must be coherent is broken down into a set of purely formal conditions¹⁸⁶ whose precise nature would take us away from value theory and too far into the theory of rational choice. Preferences are considered, says Gauthier, if and only if there is no conflict between their behavioural and attitudinal dimensions—specifically, preferences as they are revealed in choice behaviour must not conflict with the preferences as they are expressed in speech—and they are stable under experience and reflection.¹⁸⁷ Concerning the first of these requirements, Gauthier says, ‘If a person’s revealed and expressed preferences diverge, then her values are confused and she lacks an adequate basis for rational choice.’¹⁸⁸ Gauthier’s requirement concerning experience is similar to that of Sobel discussed above. The purpose of experience in this context is really about attaining a more accurate picture, as I called it, of the objects of our attitudes. As Sobel pointed out, if an attitude is not responsive to the true nature of its object, it is in a sense not really for it, and our informed attitudes are what we ordinarily would refer to as what we ‘really want’. Reflection plays a similar role to that of experience. Gauthier gives an example of a young lady who says ‘Yes’ to a proposal and only momentarily later realizes her mistake.¹⁸⁹

¹⁸⁵ Gauthier (1986) (pp. 24–25)

¹⁸⁶ *Ibid.* For ordinal preference rankings, these conditions are completeness and transitivity. Completeness requires the chooser to prefer one option over the other or else be indifferent between them. Transitivity requires (for example) that if you prefer A over B and B over C then you cannot prefer C over A. These two conditions are together sufficient to weakly order a set of objects (pp. 39–41). For interval measures, other conditions are needed (pp. 45–46).

¹⁸⁷ *Ibid.* (pp. 32–33)

¹⁸⁸ *Ibid.* (p. 28)

¹⁸⁹ *Ibid.* (p. 29)

Gauthier's diagnosis is that she fails to reflect before accepting. When she does come to reflect she realizes that her firm preference is rather expressed by 'No'. The analysis offered by Gauthier involves a distinction between tentative and firm preference. On this analysis, it is not the case that the lady misstates her firm no-preference by saying 'yes', nor do her initial yes-preference and her subsequent no-preference stand on equal footing. It is rather the case that 'without due consideration, she forms a tentative preference revealed in her acceptance, which she then revises.' Not all preferences are equal in status: '[T]he rejection of a tentative preference is not to be equated with the alteration of a firm or considered preference'.¹⁹⁰ In the light of this, perhaps we should say that a tentative attitude grounds a tentative value? The young lady's tentative yes-preference confers a tentative positive value on the proposal, whilst her firm no-preference confers a firm negative value on the proposal. Accepting the proposal was thus merely tentatively good relative to the young lady and only for a brief moment. It was firmly bad relative to her from soon thereafter onwards, unless she revises her firm no-preference. Alternatively, a subjectivist might here say that merely tentative preferences do not constitute values at all, not even of a tentative kind. Perhaps because, unlike Gauthier, it is held that tentative preferences are not really preferences at all. Whatever the correct psychological analysis of such an event, intuitively we probably all know what is meant when it is said that the lady didn't 'really want' to accept the proposal: it wasn't her genuine desire.

Why think that only preferences that are rational (in Gauthier's sense) ground values? A potential concern about this suggestion is that speaking in terms of conditions for rational preference smuggles in an objective standard for preference, and in tying value to preference we thus set up an objective standard for value. This would clearly go against the spirit of subjectivism. However, as Gauthier points out, 'we must not be misled in so speaking, for none of the conditions addresses the *content* of particular preferences. Not the particular preferences, but the manner in which they are held, and their interrelations, are the concern of reason.'¹⁹¹ In this context Gauthier expresses his agreement with Hume's notorious claim that it is 'not contrary

¹⁹⁰ Ibid (p. 30)

¹⁹¹ Ibid (p. 25)

to reason to prefer the destruction of the whole world to the scratching of my finger'.¹⁹² Considered in itself, the rationality of a particular preference cannot be assessed. It is only the manner in which it is held and its relation to other preferences that can be judged to be rational or irrational. Gauthier insists that practical rationality is strictly instrumental.¹⁹³

[I]n identifying rationality with the maximization of a measure of preference, the theory of rational choice disclaims all concern with the ends of action. Ends may be inferred from individual preferences; if the relationships among these preferences, and the manner in which they are held, satisfy the conditions of rational choice, then the theory accepts whatever ends they imply.¹⁹⁴

In the end, I believe that Gauthier's conditions are best seen as practical constraints arising not from value theory but from the theory of rational choice. The conditions concern not the nature of value as such, but rather the nature of the practice of pursuing value. For Gauthier, this involves maximizing: trying to get as much of the good as possible. As Gauthier himself points out, he is concerned 'not so much with preference itself, as with the requirements of any account identifying rationality with the maximization of the measure of some relationship among the possible outcomes of actions'¹⁹⁵, noting also that '[m]aximization itself imposes conditions on preference'.¹⁹⁶ If you are going to put your values into practice—which, given perspective internalism, you are going to be motivated to do—then you better do it in a rational manner. This, in turn, means that you better pursue values that are coherent and considered, since it is futile or impossible to realize incoherent and ill-considered values.

I believe that value subjectivists can accept these conditions for rational preference. But it is not clear that they should all be treated as idealizing conditions for attitudes to ground values. These conditions, or something

¹⁹² *Treatise* (ii.iii.iii.)

¹⁹³ Gauthier (1986) (p. 25)

¹⁹⁴ *Ibid* (p. 26)

¹⁹⁵ *Ibid* (pp. 39–40)

¹⁹⁶ *Ibid* (p. 38)

like them, may be taken to apply to valuing subjects *qua* rational agents rather than simply *qua* valuing subjects. The rationale for some of these restrictions is then going to be different from the one discussed in the previous section. The important thing for the subjectivist is that the rationale for these restrictions appeals to the nature of instrumental rationality and not to some objective standard of value.

A Restriction to Second Order Attitudes?

There might be several further ways to idealize attitudes. I do not claim that the forms of idealization that I have been discussing exhaust the possible idealizations available to subjectivists. However, I do think that further idealizations would also need to be anchored in the same types of rationale as I have appealed to above. In the next chapter I discuss a form of idealization that I think goes beyond this rationale. Presently, I wish to raise a proposal that focuses on second order attitudes, attitudes about attitudes. David Lewis defends such a proposal:

[W]e'd better not say that valuing something is just the same as desiring it. That may do for some of us: those who manage, by strength of will or by good luck, to desire exactly as they desire to desire. But not all of us are so fortunate. The thoughtful addict may desire his euphoric daze, but not value it. Even apart from all the costs and risks, he may hate himself for desiring something he values not at all. It is a desire he wants very much to be rid of. He desires his high, but he does not desire to desire it, and in fact he desires not to desire it. He does not desire an unaltered, mundane state of consciousness, but he does desire to desire it. We conclude that he does not value what he desires, but rather he values what he desires to desire.¹⁹⁷

In the sense of thinking or judging valuable, Lewis is right that 'to value' something is not the same as to desire it. But neither is, I believe, to value something in this sense the same as desiring to desire it. To bring Lewis's

¹⁹⁷ Lewis (1989) (p. 115). In a footnote, Lewis acknowledges that often in decision theory and economics value is identified with a measure of desiredness, and that all desires (first and well as second order desires) are taken to count equally. But Lewis distances himself from this by saying simply that 'it's not the sense we want here' (ibid, n. 2).

proposal into the context of my discussion, we must read it as applying to valuing in a different sense of the term—that which is relevant for grounding value.¹⁹⁸ It is quite possible to value something in this second sense without valuing it in the first. Interpreted in this way (which admittedly might not have been Lewis' intention), the proposal amounts to a restriction: it is not what is desired or otherwise favoured, but instead only what is desired to be desired that grounds and constitutes value.

Are there any merits to this proposal? Whilst there certainly is something to the idea that there is an intuitively relevant difference (relevant to value, that is) between what we desire and what we desire to desire, it is questionable whether this captures the difference between desiring and valuing. At least, I do not see that Lewis has given us enough by way of an argument in order to establish his conclusion. We might still ask on what grounds first order attitudes are excluded from grounding values. Whilst I do not find it inconceivable that a value subjectivist could be able to argue for a restriction of this kind, I do not think that the subjectivist needs to, or indeed should, accept any such restriction.

It seems to me rather that the most reasonable path for a subjectivist to take here is that of acknowledging that first order attitudes ground values as well as second order attitudes. But this does not mean that the distinction between first- and second order attitudes lacks bearing on value theory. Take Lewis' thoughtful addict. Insofar as he desires it, and insofar as he has an accurate picture of it, his euphoric daze is good relative to him. But insofar as he also has a second order desire not to have that first order desire, it is a bad thing relative to him that he has that desire. There is a difference between attitudes on these different levels, but the difference does not seem to be that second order attitudes ground values whilst first order attitudes do not. I thus see no reason for a subjectivist to exclude first order attitudes from grounding values.

Perhaps a subjectivist would want to go further and argue that second order attitudes may ground a kind of 'second order values'. On this proposal, it may be a bad thing that something is good or a good thing that something is bad! In the case of the thoughtful addict, insofar as he comes to recognise that his euphoric daze is good, *this*—and not merely that he

¹⁹⁸ In Chapter 8 I argue that Lewis's dispositional theory of value is not a form of subjectivism in my sense of the term.

desires it—can become the object of one of his negative attitudes, conferring upon it a negative value. It would now not merely be the case that having a certain desire is a bad thing, but *that the euphoric daze is good* would be a bad thing.

Restrictions to Non-Trivial, Non-Arbitrary and Non-Manipulated Attitudes?

Further proposals for restriction of value-grounding attitudes concern trivial, arbitrary and manipulated attitudes. Starting with trivial preferences, we can ask whether they ground values. What about someone's preference for tying his shoes in a particular order? Does subjectivism imply that tying the shoes in this particular order is of value relative to the person?¹⁹⁹ Subjectivists can agree that we wouldn't normally think of such a trivial thing as being a *value*. But this, the subjectivist can insist, is simply because it is a trivial value, and the reason that we wouldn't normally speak of this as a value is simply that when we think of paradigmatic examples of values we tend to think of great and important values, such as knowledge, peace, friendship, untouched wilderness, great works of art, and so on, and not trivial values.²⁰⁰

Adam D. Moore says, 'Since preferences can be arbitrary, and according to this view [Hobbes' subjectivism] value is intimately tied to preferences, this arbitrariness will contaminate value theory.'²⁰¹ How will the subjectivist handle this charge of arbitrariness? First, it should be asked against what standard the arbitrariness of an attitude is supposed to be judged and determined. If arbitrary means not reflecting objective values then it is clear that the charge begs the question against the subjectivist. The subjectivist is indeed committed to saying that all attitudes are completely arbitrary in this sense, but this seems not to be the sense that is relevant to

¹⁹⁹ I thank Andrés García for suggesting this objection as well as for the example.

²⁰⁰ Triviality might perhaps itself be seen as a value that is grounded on a second order attitude towards the first order attitude of tying one's shoes in a particular way?

²⁰¹ Moore (2004) (p. 76)

the objection, and this cannot be the sense that Moore has in mind.²⁰² Perhaps he means instead that attitudes can be directed towards things to which we do not attach great importance. If so, we again confront the previous concern about triviality. If a desire to tie ones shoes in a particular way is treated as a paradigmatic attitude, and knowledge, peace, friendship, etc., are taken to be paradigmatic values, then it is indeed hard to see the plausibility of the latter being grounded on the former. But this is not what subjectivists believe. Another possible interpretation of the charge of arbitrariness is that attitudes are somehow momentary and fleeting whilst values are consistent and eternal. It may be that some of our attitudes are momentary and fleeting, but our most deeply held attitudes such as our desire for self-preservation, the love and respect we have for our family and friends, our preferences for knowledge over ignorance and for peace over war, our appreciation of wilderness and admiration for great works of art, etc., can persist throughout a whole lifetime. Recall also that it is not only the objects of immediate conscious attention that are being conferred value by your attitudes. A great piece of art is not only valuable whilst it is being looked at.

Moore is also concerned about cases in which desires are contrived. 'We can imagine a situation where a child's preferences are manipulated so that the child prefers a particular kind of life or detests certain people.'²⁰³

One attack levelled at many religious groups is that they manipulate the affections of children to ensure room for religious belief. A brief look at the history of religion and the diversity of religious belief would indicate that such manipulation is possible. It would seem that few people would come to desire bloodletting or ritual sacrifice to appease the gods if these views were not shoved down the throats of the young.²⁰⁴

Clearly, subjectivists and objectivists can agree that manipulating a child in this way is a bad thing (bad *for* the child, or bad *period*, or probably both).

²⁰² Or can it? At the very beginning of his (ibid) paper Moore speaks of value subjectivity and Humean instrumental rationality as the 'twin diseases' that have 'infected moral theory' (p. 75).

²⁰³ Moore (2004) (p. 76)

²⁰⁴ Ibid

When we imagine this case our present, actual disapproval of such manipulation grounds its disvalue relative to us. Needless to add, it is not that we disapprove of it that makes it bad but the features of the case in virtue of which we disapprove of it. If we imagine a hypothetical situation in which our own desires would be manipulated in such a way that we would come to desire bloodletting and ritual sacrifice for its own sake, and no longer care about knowledge, peace, friendship, etc., we disavow our manipulated desires. As I have argued, when our present, actual attitudes extend to hypothetical situations in which our preferences are different from what they in fact are, it is our actual preferences that ground value in the hypothetical situations. Bloodletting and ritual sacrifice thus remain bad even in the hypothetical situation in which we have been manipulated into desiring these things for their own sakes.

But imagine instead that it is revealed to you that one of your actual, present desires is a result of manipulation. Let us say that you were hypnotized and a desire for X was instilled in you, after which you instantly forgot that you had been hypnotized. This new knowledge might cause your desire to change so that you now come to disfavour X. Perhaps the very fact that your previous preference for X was a result of manipulation is now itself one of the things for which you disfavour X, in which case this relational feature of X is now a bad-making feature relative to you. As previously noted, the subjectivist is not committed to saying that X was good relative to you before you gained knowledge of the manipulation, whilst after this moment X is instead bad. The disvalue of X is grounded on your present negative attitude towards it, and by hypothesis this attitude is not sensitive to your previous manipulated desire. If, however, this new knowledge (that one of your present attitudes is the result of manipulation) would not cause your preference to change or to go away, the subjectivist must say that X retains its positive value relative to you. But remember that it does not follow from this that X would be good *for* you.

5. WHOSE ATTITUDES? INDIVIDUALIST RELATIVISM AND UNIVERSALIST ABSOLUTISM

[S]ince different men desire and shun different things, there must needs be many things which are *good* to some and *evil* to others.

Thomas Hobbes, *De Homine*

I. Introduction

As previously argued only attitudes that someone actually has at present are thought to ground values. I also argued that the relevant attitudes should be restricted to conative attitudes: attitudes that involve a motivational component. A further question arises concerning whose attitudes should be taken to ground values. Is it 'my' actual, present, conative attitudes: those of the individual subject? Or, is it 'our' actual, present conative attitudes: those of all subjects that make up the constitutive ground of value? According to what I call individualist relativist subjectivism, the relevant attitudes are those of the individual subject. It is the individual's own attitudes that determine which things are good or bad relative to the individual in question, though her attitudes may be idealized in the modest sense argued for in the previous chapter. On universalist absolutist subjectivism, by

contrast, values are grounded on the attitudes of all actual subjects somehow taken together.

Universalist absolutist subjectivism comes in different forms depending on how one believes that the attitudes of different individuals should be ‘taken together’. One proposal that has been formulated (but not endorsed) by Johan Brännmark is that value is grounded on collective attitudes. Brännmark speaks of a ‘social constructivism’ according to which values are constituted by attitudes—hence this is a form of subjectivism in my sense of the term—and that conceives of the act of value constitution as ‘a collective one’ that ‘is the work of a plural subject’.²⁰⁵ Such a social value constructivist then takes these collectively constituted values to have normative authority over individuals quite independently of their own attitudes. There is thus a standard of value that is grounded on attitudes—and hence subjective—yet is thought to transcend and overrule the attitudes of the individual subject—and in that sense is universal and absolute. Interesting though it is, I raise this social constructivist version of subjectivism only to put it to one side. The view that I will take to represent universalist absolutist subjectivism is different in important respects, even though it does share some features with the social constructivist view just sketched.

On the view I have in mind, the attitudes on which values are thought to be grounded are those of an idealized ‘participant’: an actual, individual subject who has identified with all other actual subjects, taken over their attitudes and made them her own. No values are thought to be constituted by the individual subject until after she has gone through this process of universal attitude takeover. Through this ‘universalization mechanism’, diverse and conflicting attitudes of separate individuals are transformed into a single unique set of values relative to the actual world. Since this is a version of subjectivism, it is still true on this view that things can be good or bad only relative to a particular first-person perspective. But the universalist absolutist subjectivist holds that the correct perspective to take is not one’s own individual perspective as determined by one’s own present attitudes, but instead an idealized perspective generated by taking everyone else’s attitudes into account and treating them on a par with one’s own. It is thus

²⁰⁵ Brännmark (2002) (p. 94). In a footnote, Brännmark speaks of the possibility of ‘understanding the constitution of values as being done by a plural, rather than an individual subject’ (n. 52, p. 95).

not from the standpoint of each person taken as she is that values are thought to be constituted, but rather from the idealized standpoint of a person who has identified with every other person and made their preferences his or her own. This then yields a single unique set of values relative to this idealized standpoint which is then taken to be the same for all actual subjects regardless of their individual attitudes and in that sense universal and absolute.

The individualist relativist subjectivist, by contrast, denies that only attitudes thus universalized ground values, and holds instead that each individual subject's (not necessarily universalized) attitudes ground values relative to the individual subject in question. The attitudes of others enter the picture only insofar as that individual's attitudes are sensitive to their attitudes. On this individualist view there are thus as many value-grounding perspectives in a particular world at a particular time as there are individual subjects with attitudes in that world at that time. Each such perspective gives rise to a distinct set of values. We can still, on this view, speak of common or shared values, but only among those individuals who share the same attitudes.

In this chapter, I will further characterize these two rival versions of value subjectivism and pit them against each other. Though I ultimately leave the debate between the two versions in an unresolved state, I argue that there is a conditional presumption in favour of the individualist relativist version of subjectivism over the universalist absolutist version. That is, insofar as we are subjectivists it is natural to go with the individualist version.

II. The Spectator and the Participant

As previously noted, because the first person perspective does not give us access to objective values, Jan Österberg—the champion of the satisfaction interpretation that we first encountered in Chapter 2—urges us to move to the transcendental, third-person perspective, the perspective of a detached spectator. Ordinary persons are thought to be able to put themselves in or at least approximate this objective perspective by detaching themselves from their own personal engagements in the particular situation in question.

When viewing things from this detached perspective, Österberg insists, we can see that preference satisfaction and nothing else is finally valuable. Other objectivists claim that, from the same perspective, we can see that other things are valuable.

As an alternative to Österberg's spectator, who evaluates the situation objectively and from the outside, Wlodek Rabinowicz presents the participant, who instead views the situation from within. But not just from the individual subject's own first-person point of view, but also from the various first-person perspectives of other subjects. The way to do this, he claims, is by identifying with those others and taking over their preferences. Instead of distancing oneself from one's own particular perspective, as the objectivistic 'spectator model' advises, the 'participant model' wants each of us to identify with (all?) other subjects and make their preferences our own. Where the spectator model recommends a 'detached objectivity', the participant model recommends instead a 'universalized subjectivity'.²⁰⁶

Rabinowicz and Österberg are correct in identifying the spectator model as essentially objectivistic and the participant model as subjectivistic.²⁰⁷ However, though subjectivists are committed to the rejection of the spectator model (since they deny that values can be seen from the third-person viewpoint), they are not necessarily committed to the participant model—at least not insofar as following the procedure recommended by this model is thought to be obligatory for evaluating subjects. Holding that value is grounded on attitudes—which, as I have argued, is the defining claim of value subjectivism—does not commit one to holding that only attitudes that have been appropriately universalized in the manner prescribed by the participant model ground values.

We may distinguish a weak and a strong interpretation of the participant model. On the weak interpretation, evaluating subjects may identify with some or all others and take over their attitudes insofar as they are willing to do so: that is, insofar as their attitudes are sensitive to others' attitudes. On the strong interpretation, by contrast, the evaluating subject is thought to be required to take over the attitudes of others and to treat their attitudes as on a par with her own, regardless of whether she wants to do

²⁰⁶ Rabinowicz & Österberg (1996) (pp. 2–6)

²⁰⁷ The spectator model is associated with Kant and the participant model with Hume.

so.²⁰⁸ (I will return to the nature of this requirement later.) The individualist relativist need have no problem with the weak version of the participant model but rejects the strong version.

‘The preferences to be considered are ours, and not just mine.’ This is one of the mutual starting assumptions of Rabinowicz’ and Österberg’s discussion.²⁰⁹ But whilst this assumption may be natural with regard to a detached spectator with no attitudes of its own, an ordinary individual subject may reasonably ask ‘Why should I treat the attitudes of others as if they were my own?’ The individualist relativist concurs: why *should* the evaluating subject treat the attitudes of strangers, or those of his enemies for that matter, as being on a par with his own, or those of his loved ones? Though valuing subjects may be willing to take over the attitudes of some or all others and treat them as on a par with their own in deliberation, it may be asked on what basis they are required to do so. On what grounds can an evaluation made without first having taken over the attitudes of others be considered defective?²¹⁰

Indeed, we might ask why anybody’s attitudes—including one’s own—should be given a special status in evaluation. We have previously seen that attitudes may figure exclusively in the background and not enter the content of my evaluative deliberation. The features of valuable objects for which they are favoured by subjects need not involve attitudes. The very fact that something is favoured by me need not and generally is not itself counted among the features that make the thing in question valuable. That I

²⁰⁸ There might be a certain tension here arising from the idea that attitudes may figure only in the background and not in the foreground of our valuing. Is there an interpretation of what it involves to treat the attitudes of others as being on a par with one’s own that allows that this activity may go on entirely in the background and not enter the content of our deliberations? (More on this shortly.)

²⁰⁹ Rabinowicz & Österberg (1996) (p. 1). Rabinowicz makes the same assumption in his (1996) essay.

²¹⁰ As Gauthier suggests, one way to underpin this would be to appeal to a ‘universalistic conception of rationality’ and to argue that ‘since rationality is identified with the maximization of value, and rationality is universal, then what is maximized, value, must similarly be universal—the same from every standpoint’ (p. 52). Henry Sidgwick might perhaps be thought of as having defended a conception of rationality along these lines. But, as Gauthier points out, insofar as we remain true to the standard, ‘economic conception’ of rationality, this route to absolutism is closed off.

desire something is generally not one of the considerations that make it good in my eyes. On the participant model, the evaluating subject is supposed to first identify with other subjects and make their attitudes her own, and only then proceed to evaluate the object in question against the background of this entire set of assimilated attitudes. Whilst it is still true on this picture that the attitudes need not be seen as themselves value-making, it nonetheless seems that the participant model requires the evaluating subject to give attitudes a special status in her evaluation in a way that is not required by the individualist version of subjectivism. That attitudes enjoy a special status in the subjectivist theory is obvious, but why should attitudes be given a special status also for evaluating subjects?

III. Relativism and Absolutism

In the sense that the constitutive ground of value is always taken to be relative to some particular first-person perspective, it could be said that subjectivism is inherently relativistic.²¹¹ Whilst this is correct as far as it goes, I prefer to use the labels *relativism* and *absolutism* in a different way that allows me to acknowledge a distinction between absolutist and relativist versions of subjectivism. In this I follow David Gauthier, who says that

an absolutist may suppose that values are affectively determined [i.e. that value is subjective], so that a state of affairs is good simply because it is preferred by some person, but then [such an absolutist subjectivist] must

²¹¹ All subjectivists hold that value is relative to a world-bound perspective, and anyone who denies that value is thus relative must be an objectivist. However, not all objectivists deny that value is relative in this sense, so the relativist/absolutist distinction does not coincide with the subjectivist/objectivist distinction. Subjectivist relativists and objectivist relativists both hold that value is relative to a perspective, they differ in that the former affirm while the latter deny that this perspective is determined by attitudes.

hold that the state of affairs is good, not only in relation to or from the standpoint of that person, but from the standpoint of every person.²¹²

In Gauthier's terminology, 'an absolute or universal conception of value' is any conception on which values are thought to be 'the same for all persons, or for all sentient beings'.²¹³ Such a conception can be combined with either a subjectivist or an objectivist conception of value. If absolutism is combined with objectivism, we get a view according to which there is a standard of value that is not determined by attitudes (and hence objective), and that is thought to be the same for everyone (and hence universal and absolute). G. E. Moore could perhaps be thought of as having held a view of this type. His views will be touched upon later.

If instead an absolutist conception is combined with a subjectivist conception, we get a view according to which there is a standard of value that is determined by attitudes, but that is nonetheless taken to be the same for all persons; a standard that is grounded on attitudes, but which nonetheless is thought to stand above the attitudes of the individual subject. This subjective-yet-absolute standard must be generated in a manner that takes the attitudes of separate persons into account, but once this is done, the standard thus engendered is thought to become everyone's standard—valid not just from the perspective of those who share the relevant attitudes, but relative to each person regardless of their individual attitudes. This common standard overrules the individual's own attitudes insofar as these

²¹² Gauthier (1986) (p. 50). We should continuously be on our guard concerning the *because* in 'because it is preferred by some person'. As previously argued, this is dangerously ambiguous. What subjectivists hold, Gauthier included, is not that being preferred by someone is itself a value-making feature—something belonging to the supervenience base of value (except in those cases where a state of affairs involving some person having an attitude towards something is itself the object of the preference). Another thing to note is that taken on its own, the phrase 'from the standpoint of' might be thought to mean *according to*. But one can infer from other things Gauthier says that he did not intend it in that sense. Gauthier is not referring to the epistemic standpoint of the person, the standpoint as determined by the beliefs of the person whose standpoint it is. Rather, he had in mind the standpoint as it is determined by the preferences of the person.

²¹³ Ibid

do not mesh with the resulting universal standard.²¹⁴ The relevant value-grounding perspective—the perspective relative to which *my* values are grounded—is thus determined and fixed not by my own attitudes but by the attitudes of all. As noted, there are several distinct views of this general kind²¹⁵, but I have taken Rabinowicz’ participant model as the primary representative of universalist absolutist subjectivism.

On Rabinowicz’ view, as we will see, it is not quite the case that a state of affairs is ‘good simply because it is preferred by some person’ (as Gauthier put it), but rather that the state in question is *prima-facie* good in being preferred by some person. In order for the state to be good *all things considered*, the preference on which its value is grounded must first have passed a certain kind of universalization test: would the preference in question be retained by a subject who had, by identifying with all others, taken over their preferences and made them her own? If the preference in question would not survive this idealizing process, then the subject who holds this preference would in fact prefer a *bad* state—bad in an absolute, yet subjective sense. This badness is grounded on attitudes (and hence subjective), but is thought nonetheless to transcend and overrule the attitudes of the individual subject (and hence absolute in the present sense).

Relativist views stand opposed to absolutist views of all kinds. On relativist views, what is good or bad is not the same for everyone. Gauthier makes a distinction between individualistic and non-individualistic forms of relativism. On individualist relativist views, ‘each person has his own good (and bad)’ and ‘the goods of different persons are not parts of a single,

²¹⁴ In virtue of being subjectivists, adherents of such views must hold that things can be good or bad only relative to a perspective. But in virtue of being absolutists, they must hold that a particular perspective enjoys a privileged status in relation to other perspectives. The perspective of the idealized subject who has taken over the preferences of all others and made them his or her own is thus thought to enjoy a privileged status over the non-idealized perspective of the individual subject who has not gone through such a process of universal attitude take-over. The privileged status in question is the status of being constitutive of value.

²¹⁵ Gauthier (1986) puts John Stuart Mill in this category. Gauthier takes Mill’s claim (in Chapter IV.3. of *Utilitarianism*) that ‘the sole evidence it is possible to produce that anything is desirable is that people actually desire it’ to suggest a subjectivist understanding of value, while Mill’s insistence ‘that each person’s happiness is a good to that person, and the general happiness, therefore, a good to the aggregate of all persons’ to suggest an absolutist understanding (p. 52). Perhaps R. M. Hare (1981) is a better example than Mill.

overall good.’ Non-individualist relativists, by contrast, ‘suppose there to be a good common to a family, or to a class, or to a society, while denying that the goods of different families or classes or societies are parts of an overall good’.²¹⁶ Like Gauthier, I leave non-individualistic forms of relativism aside. Within a subjectivist framework, given the prominent role that attitudes are taken to play, individualist relativism comes very naturally, since the individual is precisely the entity that has attitudes.²¹⁷

Gauthier’s own view, as we have seen earlier, is a version of subjectivism in that it sees value as a ‘product of our affections’ and as ‘created or determined by preferences’, but unlike absolutist versions of subjectivism, Gauthier’s subjectivism is relativistic and individualistic in that it sees value as being created or determined by each individual’s own preferences. He writes ‘[W]hat is good is good ultimately because it is preferred, and it is good from the standpoint of those and only those who prefer it’.²¹⁸ There is no single, universal standard of value that is valid for every person and that may overrule the standard set by the preferences of the individual subject. Things are good or bad only relative to the perspectives of individual subjects as determined by each individual’s own preferences. The preferences of others are thought to enter the picture only insofar as the individual’s own preferences are sensitive to them.²¹⁹

As Gauthier notes, Thomas Hobbes held a view of value that combined subjectivism with individualist relativism. Consider first the following oft-quoted passage from *Leviathan*:

²¹⁶ Ibid (p. 50)

²¹⁷ I take it that in so far as groups of various kinds can be thought of as having attitudes, this is so only in a derived sense.

²¹⁸ Gauthier (1986) (p. 50). Note again that the phrase ‘because it is preferred’ should here be understood in accordance with the object interpretation. The fact that an object is being preferred is not taken to be part of what makes it valuable, instead value is taken to accrue to the object of the preference in virtue of the non-evaluative features of the objects *for which* it is being preferred.

²¹⁹ Whether our attitudes commonly are thus sensitive is an empirical matter, but I suspect that they very often are sensitive to the attitudes of some others, but rarely if ever to the attitudes of all others. It is widely acknowledged in evolutionary anthropology that humans are naturally biased in favour of their own kin and those with whom they (or may potentially) enter into reciprocal relationships.

[W]hatsoever is the object of any man's appetite or desire; that is it, which he for his part calleth *good*: and the object of his hate, and aversion, *evil*; and of his contempt, *vile* and *inconsiderable*. For these words of good, evil, and contemptible, are ever used with relation to the person who useth them: there being nothing simply and absolutely so; nor any common rule of good and evil, to be taken from the nature of the objects themselves; but from the person of the man.²²⁰

That there is no standard of good and evil to be taken from the nature of the objects themselves but only from persons is a denial of objectivism, later echoed in Hume's insistence (discussed in Chapter 2) that we can never find vice and virtue as long as we consider the object but only when we turn our reflection on ourselves. That there is nothing simply and absolutely good I take to be a denial of absolutism. The relativist subjectivist endorses both of these claims. The first claim of the quote, however, that the word *good* is always used in relation to the user, is not something to which the relativist subjectivist is necessarily committed.²²¹ In *De Homine*, Hobbes wrote that 'good and evil are correlated with desiring and shunning' and 'since different men desire and shun different things, there must needs be many things

²²⁰ *Leviathan* (i.vi.vii.)

²²¹ Depending on how this claim is understood. It could be read either as expressing the (cognitivist/descriptivist) view that evaluative judgments report the speaker's attitude, that they state claims that are true or false in relation to the speaker's psychology. Alternatively, it might be read as an endorsement of the (non-cognitivist/expressivist) view that evaluative judgments express the speaker's attitude without thereby stating claims that can be true or false. The value subjectivist, as I understand him here, is not committed to either of these views, and as I have previously insisted value subjectivism is neutral between cognitivism/descriptivism and non-cognitivism/expressivism in general. A third possible interpretation of Hobbes's claim is that when we use value words such as 'good' or 'evil' we ascribe objective and absolute values to things but since nothing is simply and absolutely good or evil, only good or evil relative to the individual person, our evaluations are mistaken. In Chapter 8, I will discuss these three views further. A fourth interpretation of what Hobbes is saying is that evaluations are always *caused* by something in or about the individual evaluator, presumably his or her desires and aversions which is compatible with all three of the above views.

which are *good* to some and *evil* to others'. He concludes that 'good is said to be relative to person, place, and time.'²²²

Like absolutism, relativism too is thought by Gauthier to come in both subjectivist and objectivist versions. Eric Mack is an example of someone who has explicitly defended the combination of an individualistic relativism (or 'agent-relativism', as he calls it) and value objectivism. Mack denies the subjectivist thesis that values are 'conferred upon otherwise valueless states of affairs . . . by our preferences or desires or commitments.'²²³ Things being good or bad is 'not simply a function of our affective states', he insists.²²⁴ Yet, Mack argues that his objectivism does not commit him to absolute (or 'agent-neutral') values, and he favours 'a construal of value as agent-relative yet objective'.²²⁵ He says that 'value is essentially value for this or that specific agent'.²²⁶ Unlike the subjectivist, who holds that the relation between the agent (the subject) and value is one of grounding or constitution, Mack tells us that he has in mind 'some further and more substantive relation, for example, its fulfilling my preference, its satisfying my desire, its being constitutive of my self-realization, in virtue of which I have reason to promote [it]'.²²⁷

It is not entirely clear what Mack's view amounts to, but it seems to include certain first order commitments in addition to its second order content. Some of the things he says might be taken to suggest a relational

²²² *De Homine* (XI.4.). The passage continues: 'What pleaseth one man now, will displease another later; and the same holds true for everyone else.'

²²³ Mack (1993) (p. 221)

²²⁴ *Ibid* (p. 220)

²²⁵ *Ibid* (p. 209)

²²⁶ *Ibid* (p. 210). I will argue below that we need to distinguish what it is for something to be good for a person from what it is to be good from the person's perspective. It is not entirely clear whether Mack is arguing for an individualistic relativist conception of value, or if he is instead arguing for the view that all goodness is goodness-for, that all values are relational.

²²⁷ *Ibid* (p. 215)

rather than a relativist conception of value.²²⁸ Despite this ambiguity, I will let Mack stand as a representative of an objectivist form of relativism.²²⁹

Insofar as subjectivist and objectivist views can be combined with either a relativist or an absolutist view in the way Gauthier proposes, we get the following four different possible positions²³⁰:

	Subjectivism:	Objectivism:
(Individualist) Relativism:	(1) Value is grounded on attitudes and relative to each individual (Hobbes/Gauthier)	(2) Value is not grounded on attitudes but relative to each individual (Mack)
(Universalist) Absolutism:	(3) Value is grounded on attitudes but not relative to each individual (Rabinowicz)	(4) Value is not grounded on attitudes and not relative to each individual (G. E. Moore/Österberg)

²²⁸ In the following chapter, I elaborate on the relational/relative distinction at length.

²²⁹ The view of Adam D. Moore (2004) involves a similar ambiguity as that of Mack. Moore approvingly quotes Mack, but calls the latter a ‘relationalist’ rather than a relativist (p. 83). One could think that Moore simply fails to recognise the distinction between relational and relative, but at one point he indicates that he sees them as different. He says that ‘philosophers who defend relational accounts of value need not make value relative to individual agents’ (p. 84). Admirably, Moore recognises that the divide between relational and non-relational is independent of that between subjectivism and objectivism (pp. 83–84). His own view is that value is objective and relational (p. 88). I think that Moore should be seen as advancing a relational account similar to that of Christine Korsgaard (which I will discuss in the next chapter) and not a relativist view.

²³⁰ Some of these views might be seen to include first order content in addition to their second order content. It might be impossible to fully separate claims of the former sort from claims of the latter sort.

Sorting further philosophers into these boxes would be a fun, but also dangerous exercise, since almost any such categorization is bound to be controversial.²³¹

Gauthier holds that, even though the two are logically independent, a relativist conception of value is the natural companion to a subjectivist conception, and that an absolute conception is the natural companion of an objectivist conception. He argues that the two other combinations—positions 2 and 3 in my table—are ‘under pressure’.²³² Below I will argue that there is a presumption in favour of 1 over 3. Mack too concedes (to Gauthier) that there are ‘plausible links’ between value relativism (which Mack defends) and value subjectivism (which he rejects).

Subjectivism is a natural expression of the more general agent-relativist idea that what is valuable is valuable in and through its relation to agents. Subjectivism specifies affective relations as those in and through which states of affairs have value. And the idea that value exists as the bestowal of affection, in turn, supports, the nonneutrality of value. This support derives from a sense that the power of any preference to confer value has a limited effective range. While my desire for [a state of affairs O] may have the power to confer value on [O] for me, it is difficult to see my desire for [O] having the power to bestow upon [O] an agent-neutral value that summons all to its service. The value that my desire confers, the reason that my preference

²³¹ Rabinowicz & Österberg (1996) claim Hume and Kant as the respective predecessors of the participant and spectator models, but it would certainly be possible to claim that Hume belongs in the first box together with Hobbes. Kant should probably go in box 4, but some scholars might want put him elsewhere.

²³² Gauthier (1986). Concerning the ‘pressure against a conception both objective and relative’ (position 2 in my table), Gauthier says: ‘The view that each individual is a member of some natural kind, and that each kind has its own characteristic perfection, quite different from that of other natural kinds, is, if not widespread in secular ethics today, yet of great historical importance. But this view has rarely, if ever, embraced a relative conception of value. For the objectivity of each characteristic perfection, its role as a norm or standard against which each individual member of the kind may be judged, has been supposed to depend on considering each perfection to be a manifestation, appropriate to its particular circumstances, of a single universal good. The seemingly relative goods of the several kinds are really facets of absolute good. The demands of objectivity thus force an apparently relative conception of value into an absolutist mould’ (p. 53).

generates, seems at most to extend to the person whose desire or preference is at work.²³³

Yet, as indicated, despite acknowledging these plausible links, Mack nonetheless goes on to reject the position that combines relativism with subjectivism, and he thinks that it is the subjectivist element of it that has to go.²³⁴

IV. Implications of Relativist and Absolutist

Subjectivism

Does Subjectivism Lead to an Absurd Conclusion?

In order to further illuminate the differences between the individualist relativist and the universalist absolutist versions of value subjectivism, consider the following objection posed by Österberg:

²³³ Mack (1993) (p. 222). A couple of things in this quote deserve comment. The first is that though Mack is right in saying that the value an attitude confers extends only to the person whose attitude it is, it is nonetheless the case that a particular subject *S*'s desire for *O* does have the power to bestow upon *O* a value that, from *S*'s point of view, 'summons all to its service'. It may thus well be that, from *S*'s perspective, everyone has reason to promote or pursue *O*. Though, from the perspective of others, there need be no reason for anyone to do so. The second is that Mack's claim that 'my desire for [a state of affairs *O*] may have the power to confer value on [*O*] for me' is somewhat ambiguous between my attitude conferring a relational value—*O* being good *for* me—and my attitude conferring a non-relational value—*O* being good *period*—relative to me. This distinction will be discussed and argued for below.

²³⁴ Mack (*ibid*) argues, first, against G. E. Moore, for the coherence of agent-relative value, then he argues against the existence of any agent-neutral values and then against the subjectivity of value—this amounts to a defence of the agent-relativity-yet-objectivity of all value. Mack's arguments deserve attention, but I will not discuss them here.

A person who subscribes to the object interpretation [that is, subjectivism] will sometimes conclude that one and the same object both is, all things considered, [finally] good and also, all things considered, [finally] bad. This absurd conclusion will be reached whenever one person [finally] desires a state-of-affairs and another person [finally] desires its non-occurrence.²³⁵

Both absolutist and relativist subjectivists can respond to this objection, but they do so in different ways. The response that Rabinowicz provides rests essentially on the universalization mechanism of the participant model. On this model, as we have seen, the individual subject is supposed to identify with other individual subjects and take over their preferences. If the preferences thus taken over are in conflict with each other, this preference conflict will carry over into the preference set of the evaluating subject. The *interpersonal* preference conflict is thus transformed into an *intrapersonal* one. Whilst the conflict still remains, it does not involve an outright contradiction. Through the preference-takeover, the evaluating subject acquires opposing *prima-facie* desires,

and he must now try to amalgamate them into one consistent [final] preference. Only when he succeeds in this enterprise (which he need not do!), can he reach an all-things-considered [final] value judgment . . . until he succeeds, he must rest content with mutually conflicting value-judgments of *prima-facie* type, founded on opposing preferences.²³⁶

The individualist relativist subjectivist will respond to Österberg's objection in a different way. Whenever one person finally desires a state of affairs and another person finally desires its non-occurrence, the relativist subjectivist will say that the state of affairs in question is finally good relative to the one person and finally bad relative to the other person. The alleged contradiction evaporates since from no person's perspective is one and the same object simultaneously all-things-considered finally good and all-things-considered finally bad. If, from any particular individual's first-person perspective, O is all-things-considered finally good, it is never the case that O is at the same

²³⁵ Rabinowicz & Österberg (1996) (p. 16)

²³⁶ *Ibid* (p. 20)

time and from that same perspective all-things-considered finally bad. The absurd conclusion is thus avoided.²³⁷

A defender of the individualist version holds that (in a case such as that described by Österberg) insofar as evaluating subjects are willing to take over the attitudes of others—even though these attitudes conflict—and to treat these attitudes as if they were their own in their deliberations, they end up with intrapersonal attitude-conflicts, and thus with a value-judgment of *prima-facie* type, just as on the participant model. But if the evaluating subjects involved are not so willing, individualist relativist subjectivism and universalist absolutist subjectivism have different upshots in such cases. The latter view holds that evaluating subjects are required to take over the opposing attitudes of others and, failing resolution of the resulting attitude-conflicts, they must rest content with mutually conflicting values of *prima-facie* type. By contrast, the individualist relativist subjectivist does not assume that evaluating subjects are required to take over the attitudes of others and thus allows subjects to ground opposing all-things-considered final values. The conflicts between subjects holding such mutually opposed final values remain interpersonal and their resolution goes beyond value theory.²³⁸ As second order value theorists—looking at things from the third-person point of view—individualist relativist subjectivists rest content with good-relative-to claims such as ‘O is all things considered finally good relative to S₁’ and ‘O is all things considered finally bad relative to S₂’. Claims of this type, I have insisted, are not evaluative claims.

²³⁷ Bykvist (1996) also recognizes that relativist subjectivism ‘does not necessarily imply the absurdity that p is better than q, *and* q is better than p. For arguably, agent-relative theories claim that the relevant “better than”-evaluation is not the two-place “p is better than q” but the three-place “p is better than q relative to agent A” (p. 3). Note, however, that ‘p is better than q relative to agent A’ is not itself an evaluation. Such relative-to claims can be true from the third-person point of view, while evaluations are only made from the first-person perspectives of individual subjects. The evaluations, as made from within the first-person perspectives of A and B are just ‘p is better than q’ and ‘q is better than p’ respectively. There need be nothing relative whatsoever in the contents of A’s and B’s respective evaluations.

²³⁸ Insofar as such conflicts manifest themselves in practice, we must turn to cooperation, compromise, or, in the worst case, combat.

Individualist Relativist Subjectivism and Final and Intrinsic Values

A consequence of the individualist relativist view is that one and the same object can be of positive value relative to one individual subject and of negative value relative to another individual, depending on their respective attitudes towards the thing in question. Moreover, not only can objects have different values relative to different subjects, but the same object can also have different types of value relative to different subjects. For example, an object can be of final value relative to one individual, but of merely instrumental value relative to another. This would be so whenever the former individual favours the object in question for its own sake, whilst the latter individual favours the object merely as a means for something else that he or she favours. Martin Harvey identifies this idea in Thomas Hobbes' theory of value. Applied to the example of justice, Harvey puts the following words into Hobbes' mouth:

[While] the objects of [our] desires constitute what is good [relative to] us, the material content of these desires remains open. As such, axiological room obtains for constructing justice as either instrumentally or intrinsically valuable (relative to the desires of the individual agent).²³⁹

Harvey here uses the term *intrinsic* instead of *final*, as he doesn't make the distinction that I make between final and intrinsic value. If we allow, as I have argued that we should, that there are final values that are not intrinsic, then we should allow also that one and the same object can be finally intrinsically valuable relative to one subject while finally extrinsically valuable relative to another. This would be so whenever the former individual favours the object for its own sake in virtue of its internal properties, whilst the latter individual favours the object for its own sake in virtue of its relational properties.

²³⁹ Harvey (2004) (p. 445). Note that Harvey clearly interprets Hobbes as adhering to the object interpretation. Value is here ascribed to the objects of our desires and not to the satisfaction of our desires. Harvey does not, however, make the distinctions between good-relative-to and good-for, which is why I have substituted 'relative to' for 'for' in the quote.

Individualist Relativist Subjectivism and Common or Shared Values

It should be pointed out that the individualist relativist version of subjectivism does not have the consequence that we cannot speak of common or shared values.²⁴⁰ We can think of a common or shared value as accruing to something insofar as we have the same or similar attitudes towards it. Martin Harvey recognizes this when he comments that on Hobbes' theory of value something may be a common good insofar as it is intersubjectively desired. Taking the example of the value of peace, Harvey writes,

'Peace' thus constitutes a 'common good' in that all inter-subjectively agree that peace ought to be sought and war avoided. In turn, the 'surface objectivity' provided by a 'common' but at depth subjectively held object of desire affords Hobbes a sufficient axiological basis upon which to construct, from a third-person perspective, an objective civil science without any untoward ontological commitments.²⁴¹

At this point it might perhaps be suggested that if something is valued by everyone, then this fact alone is enough to call it an objective value. But this is not so. As John Mackie points out, 'To say that there are objective values would not be to say merely that there are some things which are valued by everyone, nor does it entail this'. Continuing the same passage, Mackie too (like Harvey's Hobbes) acknowledges that subjectivism does not exclude intersubjective values: 'There could be agreement in valuing even if valuing is just something that people do, even if this activity is not further validated. Subjective agreement would give intersubjective values, but intersubjectivity is not objectivity.'²⁴²

²⁴⁰ This was pointed out by Hobbes in *De Homine* (XI.4.): 'There can be a common *good*, and it can rightly be said of something, *it is commonly a good*, that is, useful to many, or good for the state. At times one can also talk of a good for everyone, like health: but this way of speaking is relative; therefore one cannot speak of something as being *simply* good; since whatsoever is good, is good for someone or other.'

²⁴¹ Harvey (2004) (p. 440)

²⁴² Mackie (1977) (p. 22)

V. What Kind of Relativism Is This?

What Does ‘Good Relative To’ Mean?

Mark Schroeder poses a challenge for anyone who proposes to treat value as relative to particular agents, and to speak in terms of things being good or bad *relative to* this or that person. How are we to understand such talk? We can start by noting that Schroeder is correct in his observation that, unlike ‘good’ and ‘good for’, the expression ‘good relative to’ does not occur in everyday language. He goes on to ask whether we should treat ‘good relative to’ as a purely theoretical notion, or if instead it is a notion of which we have some pre-theoretical grasp, one that we are able to have intuitions about.²⁴³ If ‘good relative to’ is not a purely technical notion (on a par with ‘electron’ in physics), there should be some way for us to talk about what is good relative to whom in terms that are familiar to ordinary speakers of English. But, as Schroeder rightly points out, attempting to do this by understanding ‘good relative to’ in terms of ‘good for’ would be a mistake.²⁴⁴ As I will emphasize at greater length in the next chapter, ‘good relative to’ and ‘good for’ are distinct notions. Whilst ‘good for’, like ‘good period’, is an evaluative notion²⁴⁵, a notion that is used to make first order claims, ‘good relative to’ is a theoretical notion used to characterize a certain second order thesis. ‘Good relative to’ is thus not meant to capture anything in our first order, evaluative and normative thought and discourse, but instead a certain philosophical idea about the essentially relative nature of value. But even though ‘good relative to’ is something of a term of art, this does not mean that it is a purely technical notion of which we have no pre-theoretical grasp. And, as I will try to show presently, it does not mean that good-relative-to talk cannot be put in familiar terms using everyday language.

²⁴³ Schroeder (2007a) (pp. 268–271)

²⁴⁴ Ibid (pp. 272–273)

²⁴⁵ Though, as I pointed out in Chapter 1, ‘good for’ may sometimes be used in a purely descriptive way.

For Schroeder, one way in which we might make sense of good-relative-to talk in more familiar terms involves speaking about it—as I have been doing all along—in terms of different points of view or perspectives. This is how I want to read good-relative-to talk. But Schroeder claims not to understand what is meant when it is said that something is good ‘from a point of view’. He points out that talk of points of view is itself open to more than one interpretation.²⁴⁶

A Relativisation of Truth?

Some of the interpretations that Schroeder considers and rejects involve the idea that evaluative claims can be true from some points of view (whilst false from other points of view). I think it should be clear from the foregoing discussion that this is not what I ascribe to the value subjectivist.²⁴⁷ It is the nature of value and not the nature of evaluative truth that subjectivists wish to understand in relative terms; it is value and not evaluative truth that is the object of the subjectivist’s relativisation. I have, however, argued that non-evaluative claims of good-relative-to type can be true. But the truth of such claims is just the ordinary, non-relative truth. Value subjectivists are not committed to any relativism about truth.

Schroeder goes on to say ‘[T]here is only one interpretation that seems to stand on its own as helping us to understand what “good-relative-to” talk is all about’. On this interpretation, he says, such talk ‘should really be understood in terms of talk about what is good that we already understand, together with a special kind of proposition-taking connective, the *point of view* connective, which we can think of as taking propositions as objects in the way that “believes that” and “desires that” do’.²⁴⁸ Schroeder is surely onto something here, but it needs to be further elaborated.

²⁴⁶ Ibid (p. 274)

²⁴⁷ An evaluative claim may, of course, *appear* to be true to one subject whilst appearing false to another, leading one of them to form the belief that the claim in question is true and the other to form the belief that the claim is false. But this is quite another matter from holding that the claim really *is* true from one perspective whilst being false from another perspective.

²⁴⁸ Ibid (p. 275). About this latter interpretation, Schroeder comments that it ‘has led some philosophers to wrongly suspect that agent-relative value is a kind of subjective value—value that only exists from points of view and, hence, is only believed to exist’ (ibid). While it is

Good Relative To and Good According To

Saying that ‘good relative to’ means ‘good from a perspective’ might suggest that it means the same as the expression ‘good according to’. But this would not be right. What is good according to you is what you think or judge to be good. What it is to make a value judgment will be discussed in Chapter 8, but we all have at least some idea of what it is to make such judgments. Whatever it is to be good or bad from a perspective, it cannot be that. If it was, there would be no need to introduce the novel concept of ‘good relative to’, since we already have ‘good according to’ which we all understand fairly well. To distinguish ‘relative to’ from ‘according to’, we should say something like this: whilst what is good according to you is what you think or judge good, what is good relative to you is what you favour. Or, in other words, we could say that whilst what is good according to you is determined by your value judgments, what is good relative to you is determined by your pro and con attitudes.

There is, however, a potential problem with understanding the meaning of ‘good relative to’ in terms of pro and con attitudes: it seems to make it into an inherently subjectivistic notion. And, as such, it would not be open for objectivists. Rønnow-Rasmussen writes that “‘good relative to *a*’ often stands for what is good from *a*’s (subjective) perspective, which is meaningful only if one accepts the conceptual possibility of subjective goodness’.²⁴⁹ If we want to leave conceptual room for objectivist as well as subjectivist versions of value relativism—which I have assumed above that we should—then it is problematic to understand ‘good relative to’ in such a committal way. A promising way out of this problem is to let ‘good relative to’ mean simply ‘good from a perspective’, and beyond this leave it open for different accounts of what determines or fixes the relevant perspective. We can then understand subjectivist relativism in terms of attitudes, and let objectivist relativists provide their own accounts of how to fix the relevant

correct that relativism does not necessarily lead to subjectivism, relativism and subjectivism are, as I have argued, natural companions and objectivist relativism as well as subjectivist absolutism are views with internal tensions.

²⁴⁹ Rønnow-Rasmussen (2013)

perspective in some way that would not presuppose that value is grounded on attitudes.²⁵⁰

Fixing the Relevant Perspective

The perspectives that subjectivists have in mind when they speak of things as being good or bad relative to this or that perspective are the attitudinal points of view of individual subjects, as determined by their actual, present conative attitudes. However, as argued in the previous chapter, some of these attitudes need to be redirected when the objects that are being favoured by the subject in fact lack some or all of the features for which they are being favoured. In such cases, I argued, it is reasonable to think that the attitude still grounds a value relative to the subject, but that this value accrues not to the object of her misdirected attitude, but instead to some other object or objects that actually possess the features in question. Attitudes that are misdirected in this sense are thus not disqualified from grounding values but are merely redirected to those objects that actually have the relevant properties that have been conferred the status of value-makers relative to the subject. In saying that ‘X is good relative to S’s perspective’, the relevant perspective is that determined by S’s redirected rather than her misdirected attitudes. A rationale for this idealization (if it deserves to be so called) was defended in the previous chapter.

The universalist subjectivist, as opposed to the individualist subjectivist, defends also a further idealization. On the universalist view, the relevant value-grounding perspective is not determined by the attitudes of the individual subject taken as she is, but instead by the attitudes of an idealized subject who has taken over the attitudes of all other actual subjects and made them her own. If this view is correct, X may be all things considered good relative to S even though S does not have a positive attitude towards X—and indeed even when she has a negative attitude towards X—and this may be so even when S has an accurate picture of X as well as an accurate picture of the features in virtue of which she favours X, and she is perfectly instrumentally rational. This is so, we have seen, because the perspective from which S’s (as well as everybody else’s) values are constituted

²⁵⁰ I take it that Eric Mack (1993) is trying to do precisely that (among other things).

is, on this view, determined not by S's own attitudes, but by the attitudes of all actual subjects. I argue below for a presumption in favour of the individualist version of subjectivism.

On Stephen Darwall's interpretation of Hobbes (which I have returned to often by way of illustration), Hobbes holds that in desiring something 'we ascribe to it a property, that of being good'.²⁵¹ There is nothing relativistic about *that*. As we have seen earlier, Darwall ascribes to Hobbes the view that all deliberation begins in the agent's desires, not from premises about her desires. The agent reasons from a premise that she accepts in having a desire. '[T]his premise is something normative—that something would be good, that she is to or ought to do something. The agent has these normative thoughts because she has desires. They are the "appearances" of her desires'.²⁵² Further, 'Thought and discourse about good and evil encode an agent's view of things in deliberation, from the agent's perspective provided by her desires'.²⁵³

From the first-person point of view of an individual subject S who favours (has a pro attitude towards) an object O, O is good *period* (or, perhaps good *for* someone), not good-relative-to-S. It is only when we move to the third-person point of view that we will come to recognize—so the subjectivist relativist holds—that all values are relative to the respective points of view of individual subjects, as determined by their pro and con attitudes. From the third-person perspective, we can see that value-constitution is going on whenever subjects are having pro and con attitudes. We can further see what the objects of their respective such attitudes are and the non-evaluative properties for which these objects are favoured or disfavoured. But while this information allows us to derive conclusions of the good-relative-to type, it does not allow us to derive anything evaluative. From the third-person perspective, we can see that S has a pro attitude

²⁵¹ Darwall (2000). Darwall's Hobbes holds that in desiring something, 'we ascribe to it a property, that of being good (something we ought to achieve), that it does not literally have' (p. 318). Theoretical knowledge of the fact that 'there is nothing in the nature of the objects of [the agent's] desires that answer to the normative properties she attributes to them in having desires' need not undermine deliberation (p. 333).

²⁵² *Ibid.* As previously argued, it is not necessary that the subject operates with a normative premise.

²⁵³ *Ibid.* (p. 334)

towards O. From this information we can derive the non-evaluative conclusion that 'O is good relative to S'. But we cannot, from this detached perspective, derive any evaluative conclusion such as 'O is good' or 'O is good *for* S'.

What Follows?

On both the individualist and the universalist versions of subjectivism, it is the constitutive ground of value that is taken to be relative, not the products of value-constitution. The nature of value-constitution is essentially relative, but this does not mean that the constituted values are themselves relative.²⁵⁴ Subjectivism—even in its individualistic version—does not force us to evaluate anything in either explicitly or implicitly relative terms. There need be nothing either implicitly or explicitly relative in the content of our evaluations. Insofar as the favouring subject ascribes a value-property to the object of her pro attitude (which need not be the case), it will be the property of being good *simpliciter*—or, perhaps, the property of being good *for* someone—not the property of being good-relative-to-the-evaluator, or some such thing. Subjectivism, as I understand it here, does not imply that the meaning of 'X is good' is 'X is good relative to me [the speaker]', or anything along such lines.²⁵⁵ Subjectivism does therefore not lead to the pernicious conclusion that genuine disagreements in evaluative matters are impossible.

²⁵⁴ As I will argue in the next chapter, at least some values might be relational, like goodness-for; but this I take to be a different matter entirely.

²⁵⁵ Whether such a view of the meaning of evaluative pronouncements is compatible with value subjectivism is a question that I will take up in Chapter 8, but it certainly is not required by it.

VI. A Presumption in Favour of Individualist Relativism

Individualist relativist subjectivism does not set up any substantial restrictions on what may be valued. It does not, for example, force us to value 'individualism' and disvalue 'universalism'. But, as noted above, the universalist absolutist version of subjectivism does force us to evaluate in a particular way. This view leans on the participant model that requires evaluating subjects to take over the preferences of others and treat them as on par with their own. What is the nature of this requirement? It seems to be a first order, normative requirement. Insisting on such a requirement therefore seems to involve going beyond second order value theory. The relativist subjectivist might challenge the absolutist subjectivist on these grounds. Without going beyond the domains of second order value theory, there seems to be no resources available for turning the diverse and conflicting attitudes of separate individuals into universal and absolute values. This is the basis of the presumption in favour of individualist relativism.

However, this challenge assumes that the debate between universalist absolutism and individualist relativism falls wholly within the second order domain. A possible response from the absolutist subjectivist might be built on the insistence that universalist absolutism is preferable over individualist relativism precisely because it fits better with a certain favoured first order, normative theory.

The spectator and the participant models can be seen, Rabinowicz and Österberg tell us, as 'two different ways of viewing utilitarianism' or as anchored in 'two different models of utilitarian evaluation'.²⁵⁶ Their value-theoretical debate is thus deliberately placed within the framework of a particular first order, normative theory.²⁵⁷ I suspect that universalist

²⁵⁶ Rabinowicz & Österberg (1996) (p. 3)

²⁵⁷ Towards the end of their joint paper, these authors suggest that given the object interpretation preference utilitarianism turns out to be a meta-ethical position. But I doubt that this is the best way to understand this.

absolutism is going to look more attractive to those subjectivists who are independently attracted to a utilitarian ethical doctrine. Gauthier notes that '[a] position both subjectivist and absolutist seems implicit in the views of many defenders of . . . utilitarianism'.²⁵⁸

However, Rabinowicz and Österberg also rightly insist that '[t]he axiological component of [preference utilitarianism] can . . . be divorced from its normative component'.²⁵⁹ and that a preferentialist axiology can be accepted by non-utilitarians of several different persuasions.²⁶⁰ In sharp contrast to this, it has been argued that, in virtue of its relativistic nature, a subjectivist theory of value cannot be combined with a consequentialist/utilitarian normative theory. Krister Bykvist has suggested an argument to this effect.²⁶¹ He first points out that if the subjectivist theory of value is seen as relativistic 'in the sense of generating rankings of a set of outcomes that vary from one agent position to another', it might be the case that '[a] state of affairs p is better than a state of affairs q relative to a certain agent position, but q is better than p relative to another agent position'.²⁶² If the individualist relativist version of subjectivism is adopted, it will be agreed that this will indeed often be the case, since subjects have different and sometimes opposing attitudes towards the same things. Bykvist continues: 'These explicitly agent-relative evaluations are then used when we, as preference utilitarians, move from axiology to normativity. Each agent ought to perform the action which leads to the outcome that

²⁵⁸ Gauthier (1986) (p. 52). Gauthier argues that utilitarianism 'finds itself under pressure to move away from a conception of value at once subjective and absolute'. '[I]f utilitarianism remains true to its roots in the economic conception of rationality, then either subjectivism or absolutism gives way. On the one hand value may be conceived as relative, but a special form of value, moral value, is introduced, which is the measure of those considered preferences held from a standpoint specially constrained to ensure impartiality [like the standpoint of the idealised participant model]. On the other hand value may be conceived as objective, as the measure of an inherent characteristic of states of experience—enjoyment—that affords a standard or norm for preference' (ibid). (Gauthier has in mind hedonistic utilitarianism, but the point may be extended to cover also a utilitarianism coupled with the satisfaction interpretation of preferentialism.)

²⁵⁹ Rabinowicz & Österberg (1996) (p. 2)

²⁶⁰ Gauthier is a case in point.

²⁶¹ Though, he ultimately does not embrace its conclusion.

²⁶² Bykvist (1996) (p. 14)

maximizes value relative to the agent.²⁶³ Bykvist then asks whether it is possible to combine such relative evaluations with consequentialism. He notes that given an influential definition of consequentialism credited to Samuel Scheffler, the answer seems to be ‘no’. This definition runs as follows:

[Consequentialism] specifies some principle for ranking overall states of affairs from best to worst from an impersonal point of view. In other words, the rankings generated by the designated principle are not agent-relative; *they do not vary from person to person, depending on what one’s particular situation is.*²⁶⁴

If ‘an impersonal point of view’ is read as meaning the third-person point of view—the transcendental perspective of a detached spectator—then consequentialism would not be compatible with any version of subjectivism, either individualist or universalist. However, if it would be allowed that the idealized first-person point of view of the participant model—the perspective of an individual subject that has identified with all others and is now treating everyone’s preferences as if they were his own—is ‘impersonal’ enough, then the compatibility of consequentialism with the universalist absolutist version of subjectivism seems unthreatened. But is consequentialism also compatible with the individualist relativist version of subjectivism that Bykvist seems to have in mind?

I think that Bykvist is right in saying that, *given Scheffler’s definition*, consequentialism cannot be combined with individualist relativist subjectivism. But, as Bykvist indicates at the end of his essay, we should take issue with Scheffler’s definition. I think that it would be a mistake to see normative consequentialism as necessarily including or presupposing any

²⁶³ Ibid. As I am understanding relativist subjectivism here, it does not imply that there need be anything relative in the content of our evaluations. Valuing is always done from the first-person perspective of the subject and it is only from the third-person perspective that values are seen to be relative to this or that individual subject. Utilitarianism, when combined with this individualist relativist theory of value, will thus not tell agents to maximize ‘relative values’ but simply to maximize value. However, in following this same uniform normative advice, the values being pursued by agent A need not coincide with those pursued by B because what is valuable from A’s point of view need not be the same as that which is valuable relative to B (as A and B may have different attitudes).

²⁶⁴ Scheffler (1987) (p. 1) as quoted by Bykvist (1996) (p. 3). Bykvist’s emphasis.

specific theory of value. At least, it seems we would need a reason for why we must or should conceive of consequentialism in this value-theoretically loaded way, especially as the normative and value-theoretical theses are readily distinguishable and that one can clearly accept one without the other. I thus maintain that even in its individualist relativist version, value subjectivism is compatible with consequentialist normative theory, as well as with other normative theories. That a consequentialist normative theory, when combined with such an individualist subjectivist theory of value, would yield different and possibly mutually incompatible practical recommendations to different agents does nothing to threaten the compatibility of the normative theory with the value theory. Though, it might well threaten our preconceived ideas of what consequentialists prescribe in many cases.

The relativist subjectivist who adopts consequentialism might end up with something similar to what Mark Schroeder calls ‘agent-relative teleology’. Whilst ‘ordinary’ consequentialism

invokes an ordering on possible states of affairs, which it uses to induce an ordering on actions available to an agent at a time in the following way: the rank of the action is given by the rank of the resulting state of affairs,

agent-relative teleology

invokes one ordering on possible states of affairs *per agent* and ranks the actions available to each agent by the rank of the resulting state of affairs in her ranking on states of affairs.²⁶⁵

Agent-relative teleology, then, ‘appeals to different rankings on states of affairs for each agent’. ‘[T]here is no single ranking—only different rankings for each agent.’²⁶⁶

The view I have in mind may or may not be identical to that described by Schroeder, but it gives every agent the very same general practical advice: to maximize value. It does not give them the advice to maximize value-relative-to-yourself, or anything relativistic of the sort. But in practice

²⁶⁵ Schroeder (2007a) (pp. 267–268)

²⁶⁶ Ibid

different agents following this same general normative advice will—given the relativist theory of value, and given that they have different attitudes—end up maximizing different values.

Unlike that of Rabinowicz and Österberg, my project is not to find the most plausible version of (preference) utilitarianism, which is a first order, normative view, but rather to find the most plausible version of value subjectivism, which is a second order, axiological view. As I have insisted, as long as we interpret the participant model in the weak way, it involves a universalization manoeuvre that will be open not just to utilitarians. The individual subject *can* thus make use of the mechanism of the participant model—identifying with others and taking over their preferences, etc.—but insofar as second order value theory alone is concerned, the demand that we must always do so lacks foundation.

As mentioned previously, requirements of instrumental rationality may put some formal restrictions of what can be valued, and above all on which values that can be realized in action, but a requirement to take over the preferences of others and treat them as on a par with one's own does not seem to be a demand that strictly instrumental rationality alone can make. A claim to the effect that there is such a requirement is, I believe, a first order, normative claim, and one that not all value subjectivists accept. Hence, without going beyond the domains of second order value theory, and taking an evaluative or normative stance, there seems to be no access to any mechanism for turning the attitudes of separate individuals into absolute and universal values.

Had my purpose been to settle the debate between universalist absolutism and individualist relativism once and for all, more would need to be said. Here I restrict myself to noting that insofar as we remain on the second order level of inquiry, there seems to be a presumption in place in favour of the individualist relativist version of subjectivism that any defender of the universalist absolutist version would need to overcome (perhaps by appealing to some first order, normative theory, or perhaps by some other means).

VII. Is This View Objectionably Self-Regarding?

In response to an objection against his Humean theory of reasons, that this theory construes normative practical reasons as objectionably self-regarding, Mark Schroeder writes,

Most versions of the Humean Theory of Reasons don't hold that the fact that [our] desires are self-interested is an essential feature . . . Most such views hold that what is important is only that they are desires. These desires need not be self-interested; they might be completely altruistic. For all that most versions of the Humean Theory of Reasons tell us, all desires might be completely altruistic. What is important is that reasons are to be explained by interests of the self, not that they are to be explained by interests in the self. [Schroeder's own version of the Humean theory of reasons] adopts this view, and is therefore weaker than the self-interest theory. This version of the 'self-regarding' objection therefore does not get a grip.²⁶⁷

Could a similar charge be launched against the individualist relativist subjectivist? We might think that the insistence that it is 'my' attitudes rather than 'our' attitudes that forms the constitutive ground of value makes this sort of subjectivism egoistic. To dispel this worry, we might look more closely at the distinction between interests *of* the self and interests *in* the self. Gauthier makes the same distinction more explicitly. On the one hand, we have the 'interests of the self, held by oneself as subject', and on the other hand 'interests in the self, that take oneself as object'. It is the former and not the latter that are identified with the grounds of value.

This distinction gives us the conceptual tools to distinguish between other-regarding (altruistic) and self-regarding (egoistic) interests and attitudes in terms of their objects. All interests are *someone's* interests in the sense that they are held by some subject, but not all interests are self-regarding in the sense that they take oneself as object (or have as their object the well-being of the person who holds the interest in question, or the well-being of any person or sentient being). It is the object of an attitude that

²⁶⁷ Schroeder (2007b) (p. 25)

determines whether it is other-regarding or self-regarding not whose attitude it is.

On individualist relativist subjectivism, values are grounded on the attitudes of the individual subject: the interests of the self, held by oneself as subject. The objects of these attitudes, however, may, but need not, have anything to do with oneself: the attitudes may, but need not, be interests in the self, taking oneself as object. A subject's final attitudes may, but need not, be self-regarding. Plausibly, some of our attitudes are self-regarding—we all take an interest in ourselves, in our own well-being (for example). Other attitudes are other-regarding; we are by nature partly egoists and partly altruists.

Individualist relativist subjectivism says that what is good and bad relative to someone is grounded on the attitudes that are held by that someone as subject; it does not say that only attitudes with a certain content—only self-regarding ones, only those that take oneself as object—ground value. The theory does not discriminate between altruistic and egoistic attitudes and interests. Therefore, it cannot plausibly be held to be in any objectionable way a self-regarding theory, despite its insistence that it is exclusively the individual's own attitudes that ground values relative to that individual.

A Note on Hobbes and Egoism

Thomas Hobbes is often seen as the archetypal egoist. Martin Harvey argues against what he calls 'the standard reading' of Hobbes and in favour of an interpretation of Hobbesian psychology very much on the lines of what I have argued above. As I've mentioned before, on Harvey's interpretation 'Hobbes's theory of value [is] understood as formally subjective but not necessarily substantively selfish.'

[F]or Hobbes, while we always act on our desires, and the objects of these desires constitute what is good [relative to us], the material content of these desires remains open.²⁶⁸

²⁶⁸ Harvey (2004) (p. 445)

Even though we always act on our own desires, the objects of our desires need not be anything good for ourselves; what we want may well be something that is good for another. Harvey cites Hobbes' insistence that 'there being nothing simply and absolutely [good]; nor any common Rule of Good and Evil, to be taken from the nature of the objects themselves', and comments thus:

Clearly, a more subjective axiology could not possibly be wrought: 'good' and 'evil' arise through an idiosyncratic process of 'conative endowment' whereby individuals, wholly as a matter of personal preference, invest various objects of desire and aversion with either value or disvalue. Note, however, that Hobbes here makes no material claim concerning what the substantive content of our desires *must* be.²⁶⁹

Hobbes probably thought that—as a matter of contingent empirical fact—many, if not most, of our desires are self-regarding. But Harvey insists that 'nothing conceptually hangs on this'.²⁷⁰ This empirical claim is not an essential part either of Hobbes' theory of human psychology or of his subjectivist theory of value.

As Harvey notes, Hobbes' notorious claim 'Of the voluntary acts of everyman, the object is some good to himself'²⁷¹ on the face of it seems to express a clear commitment to 'an egoistic psychology of the most pernicious sort'²⁷² However, Harvey proposes that what Hobbes is saying here is merely that whenever persons act of their own volition, it is always the case that they act on their own desires. Hobbes is not saying that those desires are necessarily self-regarding. Harvey continues:

All goods for Hobbes are personal goods since all goods are objects of individual desire(s). But just because something is necessarily a personal good [in this sense] does not require, with like necessity, that it be a self-regarding good. . . . my personal good, that is the object of my individuated

²⁶⁹ Ibid (p. 446)

²⁷⁰ Ibid

²⁷¹ *Leviathan* (i.xiv.viii)

²⁷² Harvey (2004) (p. 446)

sui generis desire, may, on logical grounds, just as likely be the satisfaction of the desire(s) of another.²⁷³

I believe that we could make the same point using the distinction between ‘good relative to’ and ‘good for’. The Hobbesian dictum could then be restated as follows: Of the voluntary acts of every man, the object is always something that is good relative to himself, but not necessarily good *for* himself (it could be something that is good for another, or just good *period* without being good for anyone in particular²⁷⁴). On my view, this restated Hobbesian dictum does not hold without qualification, but would need to be qualified in the light of what I argued in the previous chapter. However, it seems to me a plausible claim in its own right, and more to the present point, it surely cannot so easily be dismissed as crude egoism.

²⁷³ Ibid. Cf. Bernard Gert’s (1991) treatment in the introduction to Hobbes’ *Man and Citizen*. Gert notes that Hobbes’s claims that we always act on our own desires and that every man seeks what is good to himself, ‘do not, in fact, impose any limits on the desires of men or on what they consider to be good . . . It is perfectly compatible with [this] that a man desire to help another, that is, to act from charity or benevolence’ (p. 7).

²⁷⁴ I will argue in the next chapter that Hobbes’ claim that ‘there being nothing simply and absolutely [good]’ need not be understood to exclude that things may be good *period* (as opposed to good *for* this or that particular person).

6. OH MY GOODNESS! SUBJECTIVISM AND RELATIONAL VALUE

In what sense can a thing be good *for me*? It is obvious, if we reflect, that the only thing which can belong to me, which can be *mine*, is something which is good, and not the fact that it is good. When therefore, I talk of anything I get as ‘my own good’, I must mean either that the thing I get is good, or that my possessing it is good. In both cases it is only the thing or the possession of it which is *mine*, and not *the goodness* of that thing or that possession. There is no longer any meaning in attaching the ‘my’ to our predicate, and saying: The possession of this thing *by me* is *my* good. Even if we interpret this by ‘My possession of this is what *I* think good’, the same still holds: for *what* I think is that my possession of it is good *simply*; and, if I think rightly, then the truth is that my possession of it *is* good simply—not, in any sense, *my* good; and, if I think wrongly, it is not good at all.

G. E. Moore, *Principia Ethica*

I. Introduction

As Gauthier recognizes, ‘on both the absolute and the relative conceptions it is possible to distinguish what is good for some person from what is

straightforwardly good'.²⁷⁵ Rønnow-Rasmussen similarly observes 'that something is good relative to *a* leaves open whether the goodness in question is good, period; or good-for; or some other kind of goodness'.²⁷⁶ The idea that things can be good only relative to the perspectives of subjects—which I elaborated on in the foregoing chapter—should thus not be confused with the fundamentally different idea that all goodness is goodness-for, or as in any way being in conflict with the idea that things can be straightforwardly good or good *simpliciter*.

Connie S. Rosati seems to conflate good-from-a-perspective with good-for when she asks,

After all, what could be more obviously true than that things can be good or bad, not merely in an absolute sense, but *for us*? What could be clearer than that our lives can go better or worse, not just, to borrow Henry Sidgwick's memorable phrase, from the point of view of the universe, but *from our point of view*?²⁷⁷

Later on in the same paper, however, Rosati recognizes a distinction between relative and relational. She says that the notion of good-for 'does not involve a relativisation of goodness or normativity. Good for P is not goodness relative to P but a distinct relational value'.²⁷⁸ This is indeed what I will be arguing here. What is good from someone's point of view or relative to a person's perspective will here be treated as fundamentally different from what is good *for* a person.

²⁷⁵ Gauthier (1986) (p. 50)

²⁷⁶ Rønnow-Rasmussen (2013)

²⁷⁷ Rosati (2008) (p. 317)

²⁷⁸ Ibid (p. 330). As I argued in the foregoing chapter, the central subjectivist claim that values are grounded on attitudes does not involve a relativisation of goodness or normativity as such—thus forcing subjects to value things in ways that explicitly or implicitly relate to themselves—but instead a relativisation of the constitutive ground of value. The idea that things are good or bad from the subject's first-person point of view need not form any part of the contents of her values. 'Good relative to P' does not capture anything in our first order, evaluative thought and discourse, but instead a certain second order idea about the relative nature of value. 'Good for', by contrast, is meant to capture something in our first order, evaluative thought and discourse.

Goodness-for is an example of a relational value, which should be contrasted with straightforward goodness, or goodness *simpliciter*, which is a non-relational value. Recognizing both relational and non-relational values is something that both subjectivists and objectivists can do. One might perhaps be forgiven for thinking that the very idea that things can be good *period* would somehow be inimical to subjectivism, especially so to relativist versions of subjectivism. But just as a defender of the subjectivist conception of value outlined in previous chapters is not disqualified from analysing what it is for something to be finally or intrinsically good, neither is the subjectivist disqualified from analysing what it is for something to be straightforwardly, or non-relationally, good.

In the present chapter, I argue that holding that value is relative to the perspectives of subjects—that things are good only from particular points of view—does nothing to threaten the idea that things can be good, period, relative to particular perspectives. The thesis that the ground of value is relative should not be confused with the fundamentally different thesis that all values are relational, and neither does a commitment to the former thesis imply a commitment to the latter. Indeed, as I will explain, subjectivists will most naturally reject the thesis that all values are relational.

II. Goodness-Type Monism and Goodness-Type Dualism

In Chapter 1 I introduced a distinction between goodness-type monism and goodness-type dualism. According to the classification proposed by Rønnow-Rasmussen, goodness-type monism comes in two varieties. Moorean monists claim that all values are non-relational; all goodness is straightforward goodness, or can be fully understood in terms of, or reduced to, straightforward goodness. Hobbesian monists claim that all values are relational; all goodness is goodness-for, or can be reduced to goodness-for. As I will argue below, however, it is not obvious that Hobbes himself held a view of this type. Goodness-type dualists recognize both relational and non-relational values and claim that neither of these types of value can be fully

understood in terms of the other.²⁷⁹ The divide between these forms of monism and dualism is independent of that between subjectivism and objectivism, such that dualism and both forms of monism are (at least in principle) available to subjectivists and objectivists alike.²⁸⁰

Moorean monism

G. E. Moore held that all goodness is fundamentally of the straightforward type. In the quote from *Principia Ethica* with which I opened this chapter, Moore expressed his inability to even understand how it could be otherwise. He argued that the only ways in which we can possibly understand what it is for something to be good for a person is in terms of (1) it being simply or straightforwardly good that this person has or gets this thing, or (2) that the thing that the person has or gets is straightforwardly good (whether or not it is good *that* she has or gets it), or (3) what this person believes to be good in one of the first two senses. Further on in the same passage Moore adds that the straightforwardly good thing must belong exclusively to the person.

In short, when I talk of a thing as ‘my own good’ all that I can mean is that something which will be exclusively mine, as my own pleasure is mine (whatever be the various senses of this relation denoted by ‘possession’), is also *good absolutely*; or rather that my possession of it is *good absolutely*. The *good* of it can in no possible sense be ‘private’ or belong to me; any more than a thing can *exist* privately or *for* one person only.²⁸¹

Here it should be noted that what Moore means by the term *absolute* is rather what I am calling *non-relational* and not what in the previous chapter I have called absolute. Whilst Moore didn’t explicitly operate with a distinction between on the one hand relational/non-relational, and on the other hand relative/absolute, it should be clear that what he was attacking

²⁷⁹ Rønnow-Rasmussen (2013)

²⁸⁰ Note that goodness type monism should not be confused with either axiological monism (the first order view that there is only one final value) or with value-bearer monism (the view that all bearers of final value belong to one and the same ontological category).

²⁸¹ Moore (1993) (p. 150)

here is the idea of relational value (specifically of goodness-for) and not the idea that things are good only relative to subjects. Certainly, Moore would have been hostile also to the latter idea²⁸², but irrespective of Moore's actual views, the 'Moorean monists', as I am thinking of them here, advance a view according to which all values are non-relational—or at least that all values can be understood in terms of, or be reduced to, non-relational values, like straightforward goodness—but is silent with regard to the debates between absolutism and relativism and between objectivism and subjectivism. The Moorean monist position, as I understand it here, can be accepted by subjectivists (of both the individualist relativist and the universalist absolutist kind) as well as objectivists.

Hobbesian monism

The name 'Hobbesian monism' was inspired by Thomas Hobbes' insistence that 'one cannot speak of something as being *simply good*; since whatsoever is good, is good for someone or other'.²⁸³ However, it is not obvious that Hobbes himself was a Hobbesian monist. He might well have been a goodness-type dualist or even a Moorean monist! On a different possible reading of Hobbes, what he is expressing here is not the idea that all values are relational—that all goodness is goodness-for—but rather the idea that all values are relative to individual subjects. Saying the latter is fully compatible with any of the three positions on the 'relationality' of values: goodness type dualism, Moorean monism or Hobbesian monism. Holding that all values are relative to subjects' perspectives is compatible with saying that relative to a particular subject's perspective there can be both relational and non-relational values, or that there are only non-relational values, or that there are only relational values. Insofar as my alternative reading of Hobbes is correct, he was giving voice to what I have been calling individualist relativism rather than to what Rønnow-Rasmussen calls Hobbesian monism.

²⁸² Moore (ibid) did tie his view on good-for up with his objectivism and his universalist absolutism: "My own good" only denotes some event affecting me, which is good absolutely and objectively; it is the thing, and not its goodness, which is *mine*; everything must be either "a part of universal good" or else not good at all' (p. 170).

²⁸³ *De Homine* (XI.4.)

On this interpretation, what Hobbes was trying to say when he wrote that ‘whatsoever is good, is good for someone or other’ was not that all goodness is relational but rather that nothing can be good from a third-person point of view, but only from the first-person point of view of someone or other. Immediately preceding this claim, Hobbes insisted that good-for talk is relative, which may be taken to indicate that he recognized something like the distinction between good-for and good-relative-to; between relationality and relativity.

At times one can also talk of a good for everyone, like health: but this way of speaking is relative; therefore one cannot speak of something as being *simply good*; since whatsoever is good, is good for someone or other.²⁸⁴

My aim here is not exegetic, and I do not wish to anachronistically ascribe to Hobbes the modern distinctions with which I operate. But I do wish to insist on keeping the two views separate. That is, Hobbesian monism (which may or may not have been accepted by Hobbes himself) should not be confused with the individualist relativist version of subjectivism (which, I believe, can be more safely ascribed to Hobbes). These are positions in two separate debates. Whilst subjectivism, of both the relativist and the absolutist varieties, is compatible with Hobbesian monism, it is equally compatible with Moorean monism or goodness-type dualism. As I will explain later, far from being the case that a subjectivist must be a Hobbesian monist, it might even be the case that subjectivism is actually biased *against* Hobbesian monism.

Potentially better examples of Hobbesian monists than Hobbes himself include Richard Kraut and Christine Korsgaard. Kraut writes,

I am ‘against absolute goodness’ not because I think we should shun it and seek absolute badness instead, but because I doubt that the property of absolute goodness, as Moore understands it, has a useful role to play in moral philosophy or in everyday practical thinking. We should take seriously the possibility that absolute goodness is not a reason-giving property. There are, of course, many good things: books we should read because they are good novels, cameras with which we should take photographs because they

²⁸⁴ Ibid

are good cameras, friends we should love because they are good friends. Furthermore, that something is good *for* you, or *for* someone else, is, of course, often an excellent reason for you to value it. But are there things you should value because they are, quite simply, good? I doubt it.²⁸⁵

My discussion will focus on Korsgaard. She has recently defended the thesis that all values are relational, or, as she puts it, ‘the essentially relational nature of the good’:

I think that the notion of ‘good-for’ is the prior notion [that is, prior to the notion of ‘good’]. Or, to put it a better way, I think there is something essentially relational about the notion of the good itself. . . . I think there is such a thing as the good, only because there are creatures *for* whom things can be good; that is, creatures who can welcome or reject the things that they experience. In fact, I think that the idea of something’s being good without its being good-for someone should be rejected as unintelligible.²⁸⁶

Perhaps we should distinguish two different but related ideas being expressed here. The first is that the notion of good-for is prior to the notion of good; that goodness itself is essentially relational.²⁸⁷ It is this that makes Korsgaard a Hobbesian monist. The following two claims testify to her allegiance to Hobbesian monism: ‘[G]ood in the final sense is a relational notion—a form of good-for’²⁸⁸ and ‘All final goodness is essentially goodness-for the being whose final good it is’.²⁸⁹

²⁸⁵ Kraut (2011) (p. 7). The reason Kraut gives for his doubt is the following: ‘That is because I doubt that goodness can ever be, all by itself, a reason for valuing something’ (ibid). But that goodness is not itself a reason to value something can be held also by a Moorean monist. It need not be the goodness itself, but the non-evaluative properties upon which value supervenes that are reason-giving.

²⁸⁶ Korsgaard (2013) (pp. 4–5)

²⁸⁷ Insofar as these two claims are separate, we could say that the first is about the notions *good* and *good-for* while the second is about the (relational) nature of goodness. Here I have put these two claims together.

²⁸⁸ Ibid (p. 9). It is not clear what sense of ‘prior’ Korsgaard has in mind. At one point she argues for this priority in epistemological terms. She writes that: ‘It is not true that we need to know what is good before we know what is good-for someone, since despite its surface grammar, the notion of *good-for someone* is in fact the prior and more fundamental notion’ (p. 7). It is far from obvious that Korsgaard’s opponents on this point (the Moorean monists

The second idea being expressed in the block quote from Korsgaard is that things can be good only by being good *for* someone; nothing can be good without being good-for. Rabinowicz and Österberg call this view axiological individualism, which they characterize as ‘the view according to which nothing can be valuable without being valuable for someone (for some “individual”) . . . nothing can be good if it makes no one better off.’²⁹⁰ These authors relate axiological individualism to their debate between the object and satisfaction interpretations of preferentialism. Like hedonists, ‘who attribute value to pleasure or happiness only because they consider such things to be of value for their recipients’²⁹¹, so too the satisfaction interpretation seems to imply that ‘[i]t is good that people get what they want, because it is good for each person to get what he or she wants’.²⁹² Rabinowicz criticizes the satisfaction interpretation for clashing with the common sense view that ‘there can exist lots of things that are [finally] valuable without being [finally] valuable for any person or even for any particular creature’. Examples given include diversity of animal and plant

and the goodness-type dualists) are committed to the epistemological thesis that in order to *know* what is good-for someone, we must first know what is good period.

²⁸⁹ Ibid (p. 22)

²⁹⁰ Rabinowicz & Österberg (1996) (p. 11). Axiological individualism should not be confused with what I have been calling individualist relativism.

²⁹¹ Ibid. Interestingly, Korsgaard (2013) acknowledges the common ground between herself and the hedonist on this point. In a section of her paper entitled ‘How Hedonistic Utilitarianism was Almost True’, she writes that ‘I believe that there are good and bad states of affairs because there exist in the world beings *for* whom things can be good or bad in a specific way. The beings in question are the ones who are sentient, or conscious—roughly speaking, the animals. The thesis is unsurprising because it is a thesis that is also held by some philosophers who defend a very different philosophical outlook from my own—namely, the hedonistic utilitarians’ (p. 13). The thesis may be unsurprising, and find broad allegiance, but it is far from uncontroversial. I believe that the initial plausibility of axiological individualism is parasitic on the (overwhelming) plausibility of the much weaker thesis that value has *something* to do with conscious beings. Of course, value subjectivism captures this latter idea without commitment to axiological individualism.

²⁹² As these authors point out, however, a defender of the satisfaction interpretation might be able to avoid commitment to axiological individualism. See Rabinowicz & Österberg (1996) (n. 8, p. 11).

species, multiplicity of different ways of life and beautiful landscapes.²⁹³ The object interpretation is more in line with common sense in this regard. On the object view, as we have seen, it is the objects of our final preferences that carry final value, and this is so even when they are not finally valuable *for* anyone. The object interpretation, which I have identified with subjectivism, thus naturally leads to the common sense view that lots of things are good that are not good for anyone. That things can be good only relative to someone is, as I have insisted, a fundamentally different matter.

On Korsgaard's view, there must always be someone involved in the object of value, some person or animal for whom things can be good or bad, someone who can be made better or worse off in some way. This sets up substantial, first order restrictions on what can be valued.²⁹⁴ The value subjectivist, by contrast, holds the fundamentally different thesis that values are always relative to the first person point of view of the individual subject. On this picture, there need not be anyone involved in the object of value for whom it is good, and as such this raises no substantial restrictions on what can be valued.

Given the empirical assumption that some of our final attitudes take as their objects things that are not good for anyone in particular (things that make nobody in particular better off), value subjectivism naturally leads to a rejection of axiological individualism. However, if objects that don't involve being good for anyone are, as a matter of empirical fact, never (or for some reason cannot be) objects of pro or con attitudes, then subjectivism would imply that only things that are good-for someone are good.

A final comment on Korsgaard: When she says that 'the final good is something essentially relational, because it exists in relation to the consciousness, or the point of view, of a conscious being'²⁹⁵, she conflates

²⁹³ Ibid (p. 11)

²⁹⁴ Against the hedonists, Korsgaard insists (correctly, I think) that it is not only mere states of consciousness that can be good or bad for us, but also objective conditions. She identifies the relevant 'objects of consciousness' as the external conditions that 'bear on an animal's own condition' (p. 21). But surely not every object of my consciousness bears on my own condition. Korsgaard fails to explain why external conditions that do not bear on my own condition cannot be objects of my consciousness and thus be bearers of final value. Why cannot external conditions that bear exclusively on the condition of others or of no one in particular enter as objects of my consciousness and thus be valuable relative to me?

²⁹⁵ Korsgaard (2013) (pp. 21–22)

the distinction between, on the one hand, the claim that all values are relational, that all goodness is goodness-for (Hobbesian goodness-type monism), and, on the other hand, the claim that value is relative to first-person points of view, that things are good and bad only from the perspectives of subjects. Value subjectivists hold the latter, but they need not (and probably do not) accept the former; by no means does the former follow from the latter. If S desires O for O's own sake (and not for someone's sake in particular) then O is finally (non-relationally) valuable relative to S. It does not follow from this that O is good *for* S, or for any other person or conscious being.

In response to Korsgaard's claim that 'the idea of something's being good without its being good-for someone should be rejected as unintelligible', the subjectivist may reasonably insist that it is not the idea of something being good without being good for anyone that is unintelligible, but rather the idea that things are good without being good relative to the various points of view of subjects. Relational values such as goodness-for are no less mysterious than non-relational values such as goodness *simpliciter* as long as they are held to be not grounded on attitudes and not relative to the perspectives of subjects. Again, it is not the idea that things can be good without being good for anyone that should be considered problematic but rather the idea that values (whether relational or non-relational) can be objective and absolute.

I strongly suspect that the initial plausibility of Korsgaard's claim that the idea of things being good without being good for someone should be rejected as unintelligible is entirely parasitic on the overwhelming plausibility of the much more general idea that value must on some level have something to do with conscious beings. Once it is understood that acknowledging that things can be good without being good for anyone does not require a rejection of the more general idea, the attraction of Korsgaard's claim evaporates.

III. A Subjectivist Analysis of Goodness-For

Insofar as subjectivists and objectivists are willing to recognize relational values—that is, insofar as they are not Moorean monists—they will have to

offer analyses of such values. Insofar as they recognize both relational and non-relational values—that is, insofar as they are goodness-type dualists—they will have to be able to provide analyses of both the nature of straightforward goodness and the nature of goodness-for. A subjectivist who is also a dualist must hold that both relational and non-relational values are grounded on attitudes. Both straightforward goodness and goodness-for must be treated as products of our affections. Our attitudes must be seen as having the power to constitute both goodness and goodness-for.²⁹⁶

How, more precisely, could a subjectivist understand what it is for something to be good for someone? Gauthier suggests that ‘[a] state of affairs is good for someone if it contributes to his well-being’.²⁹⁷ But it may be questioned whether this is satisfactory as an analysis, as long as we don’t know how to independently analyse well-being in terms other than those of good-for. Moreover, as Rønnow-Rasmussen has argued, it is hard to see how goodness-for can be understood in terms of welfare, if welfare is treated as a thick value notion and good-for as a thin value notion.²⁹⁸ If both good-for and welfare are treated as thin evaluative notions, then understanding the one in terms of the other looks more promising, but is it really plausible to treat welfare as a thin notion? This may well be doubted. Further, it is not unreasonable to insist that a person’s welfare is just one out of the many different things that can be good for a person. Rønnow-Rasmussen writes ‘Despite the general agreement that well-being is good for us, an issue remains whether well-being and its constituents exhaust, analytically, what is good for a person.’²⁹⁹

As we have seen, subjectivist analyses of final value involve final attitudes. An object is taken to have positive final value insofar as it is favoured for its own sake; insofar as it is the object of a final pro attitude. Can this analysis be extended to cover relational values? One promising way

²⁹⁶ Rønnow-Rasmussen (2013)

²⁹⁷ Gauthier (1986) (p. 50)

²⁹⁸ Rønnow-Rasmussen (2011) (pp. 106–107)

²⁹⁹ Rønnow-Rasmussen (2013). As previously noted, saying that a person’s welfare might not be good for him or her does sound odd. But this oddity might be due merely to the fact that it militates against our first order intuitions about what is good for people, not due to it being conceptually confused. Welfare will be further discussed below.

to do this is to appeal to so-called ‘for someone’s sake’ attitudes.³⁰⁰ One suggestion along those lines would be the following:

An object, O, is good for a person, P, relative to a subject, S, insofar as S favours O *for P’s sake*.

Like the subjectivist analysis of final straightforward goodness introduced earlier, this analysis of goodness-for could be subjected to the proviso that S must have an accurate picture of O in order for S’s pro attitude to bestow this relational value on O. That is, O must really possess the features for which it is being favoured by S for P’s sake; O must really possess the properties that are good-for-making relative to S. Perhaps it should be added also that S must have an accurate picture of the person P as well?

An objection against this analysis of goodness-for is that we might want to say, in at least some cases, that an object that is being favoured for someone’s sake is nonetheless not good for that someone. Imagine, for example, that you desire something for the sake of a deceased friend. You might, so the objection goes, wish for your deceased friend’s sake that his or her work will become influential, or that this friend will be cleared of some criminal allegation, whilst at the same time deny that these things would be good for your friend, that they would make him or her better off. Surely nothing whatever can be good for someone who is not alive, and yet it is quite possible to desire something for a dead person’s sake. Hence, the proposed subjectivist analysis of goodness-for cannot be correct.³⁰¹

I, however, do not find this counterexample compelling. It is far from obvious that things cannot be good or bad for the dead. At least, this is not something that I think we can legitimately exclude on the second order level of value analysis. Many people have thought so³⁰², and even if they are wrong, there is no obvious reason for holding them to be somehow conceptually or metaphysically confused. One thing that might tempt us towards insisting that nothing can be good or bad for the dead is the aforementioned (unsuccessful) identification of goodness-for and well-being. If instead it would be accepted that well-being is only one out of many different things that might be good for a person, we can say that even

³⁰⁰ See Chapter 5 of Rønnow-Rasmussen (2011).

³⁰¹ I thank Johan Brännmark for raising this objection.

³⁰² Including Aristotle(?)

though it is not possible to enhance the well-being of a dead person, things other than well-being can still be good or bad for a dead person. What those other things are is a first order question, a question that can only be answered by taking an evaluative or normative stance, but two possible candidates are having his or her work be admired or having his or her good reputation restored after a mistaken allegation. So perhaps it would be conceptually or metaphysically misguided to want to enhance the *well-being* of the dead, but to wish other things for their sakes seems to me to be quite another matter. There might be some people who never have for-someone's-sake attitudes in relational to the dead. On the proposed analysis, nothing is good or bad for the dead relative to such people.

Leaving aside the example of the dead, it might still be insisted that it is quite possible to favour something for someone's sake whilst at the same time denying that this thing is good for that someone. Whilst this cannot be ruled out once and for all, the objection can perhaps be defused by somewhat dogmatically insisting that whichever example along these lines we might be able to come up with, once we reach the desired level of detail, it will turn out that the attitude involved is not a for-someone's-sake attitude but rather some other type of pro attitude, such as respect or honouring. You may well favour something in honour of a person without favouring it for that person's sake.

IV. Implications of the Analysis

I have argued that the claim that things can be good or bad only relative to the perspective of the individual subject should not be conflated with the claim that all goodness is relational; that things can be good only *for* this or that person (Hobbesian goodness-type monism). The former should also not be conflated with any first order claim about what is good for people (or non-human animals, plants, etc.). What is good relative to a person and what is good for that person are different concepts that must be kept distinct.

If a person desires or prefers something for its own sake (even when having an accurate picture of it), it does not follow that this thing is good *for* this person. What a person prefers for its own sake and what makes this

person better off need not be the same things. If I prefer for its own sake that some other person's life go well for him or her, then that person's well-being is finally, non-relationally good relative to me. This does not require that the well-being of this other person make *me* better off. Gauthier writes similarly that

contributing to someone's well-being is a source of positive value [relative to] others if it enters into their good [in my terminology: if it is good relative to them], as it does on our particular [subjectivist] relativistic position if those others prefer such a contribution. A state of affairs that is good for a single individual is then [straightforwardly] good from the standpoint of each person who prefers or would prefer the enhancement of that individual's well-being.³⁰³

If I prefer the survival of the rainforest for its own sake, then the survival of the rainforest is finally good relative to me. This does not require that the survival of the rainforest make anyone better off. The objects of S's final preferences need not be things that are good for S or involve the well-being of S, and neither need they be things that are good for anyone else or involve the well-being of any person or sentient creature. S herself (or any other person or sentient creature) need not be involved in the object of her own attitude in any way. The objects of S's final attitudes may well be things such as rainforests and great works of art.

The subjectivist who recognizes both relational and non-relational values holds that things can be finally good period without being finally good for anyone. In order for something to be finally valuable, on this view, it need not make anyone better off, or enter into anyone's well-being. There indeed seem to be many things besides well-being that can be objects of final attitudes and that might thus have final value. Though, in order to be finally valuable relative to an individual subject, the thing must enter as an intentional object of one of her final attitudes.

³⁰³ Gauthier (1986) (pp. 50–51). On Gauthier's view, a state of affairs is good for someone if it contributes to his or her well-being. As noted above, however, this identification can be called into question: it is not necessary to identify what is good for a person with that person's well-being. This depends on whether our chosen analysis of good-for is also plausible as an analysis of well-being.

Note on Welfarism and the Charge of Fetishism

It is plausible to assume that any person's own well-being is part of what is good relative to that person, since we all care about our own well-being. But well-being does not plausibly exhaust what is good relative to us, since most of us care about many other things for their own sake. We want many different things in addition to well-being both for our own sakes and for the sakes of others.

Simon Keller, in discussing welfarism, says that the welfarist's greatest weapon is the charge of fetishism. He writes that

the welfarist's accusation is that non-welfarist views take things whose value can usually be explained in terms of their tendency to make lives go better—things like resources, opportunities, freedom from interference, life, and great works of art—and then fixate upon them, valuing them even when their connection with individual interest is broken.³⁰⁴

The things listed are indeed plausible candidates for being objects of final value. Could the charge of fetishism do damage to subjectivism in virtue of the latter's pluralistic first order implications? I doubt it, because subjectivism respects, on the second order level, the values 'connection with individual interest'. If I desire or prefer, or have some other relevant final pro attitude toward, things such as resources, opportunities, freedom from interference, life, and great works of art, etc., then those things enter into my individual interests in the sense of being objects of my final attitudes. These things might not always 'make lives go better' in the sense that they enhance my own or anyone else's well-being, but I certainly do take great interest in these things. Subjectivism can thus safely ascribe final value to things like these whilst at the same time preserve the connection between value and individual interest. Recall the distinction between interests *of* the self, held by oneself as subject and interests *in* the self that take oneself as object. Subjectivists hold that value is grounded on the former. The commitment to individual interest in that sense—as the constitutive ground of value—does not stop subjectivists from ascribing final value to objects far beyond narrow well-being.

³⁰⁴ Keller (2009) (p. 91)

A charge of fetishism may actually be more readily applicable to welfarists themselves! Keller's charge can easily be turned around: welfare, or what makes people better off, is just one of the many different things that we take an interest in—just one out of the many different things that we desire for its own sake or for someone's sake. But welfarists insist that among the many diverse things we actually take interest in, only one—well-being—is really worthy of being valued. Welfarists refuse to ascribe final value to most of the things that we hold dear, although the values of these things can be explained in terms of their entering into our interests as object of our final attitudes. Welfarists refuse to ascribe value to these things even though their connection with individual interest is not broken and they fixate upon welfare alone. Is not that an example of fetishism?

7. PROJECTING VALUE: MAKING GOOD ON A METAPHOR

[T]he distinct boundaries and offices of reason and of taste are easily ascertained. The former conveys the knowledge of truth and falsehood: the latter gives the sentiment of beauty and deformity, vice and virtue. The one discovers objects as they really stand in nature, without addition or diminution: the other has a productive faculty, and gilding or staining all natural objects with the colours, borrowed from internal sentiment, raises in a manner a new creation.

David Hume, *An Enquiry Concerning the Principles of Morals*

I. Introduction

I have delineated a subjectivist view on the nature of value on which value is seen as being constituted by subjects and grounded in pro- and con-attitudes. On this view, I have previously noted, value is not understood as something that is part of the mind-independent world, but rather something that we project upon the world, as well as on the whole range of possible worlds. Value is seen as coming into being by being conferred upon otherwise valueless objects by attitudes that are directed towards the objects in question, and the resultant values are taken to supervene on precisely those non-evaluative properties of the objects for which these objects are being favoured or disfavoured. But how, more precisely, are we to

understand this bestowing or conferring of values? Does using the notion of *projection* add something to the subjectivist picture beyond what has been described in the preceding chapters? Rønnow-Rasmussen suggests that though many subjectivists probably are projectivists, the subjectivist is not necessarily *committed* to projectivism. ‘Establishing a relation between an object and a final attitude is what is needed (perhaps all that is needed) to generate an act of [value] constitution’.³⁰⁵ But it is unclear what more than this is thought to be involved in a commitment to projectivism.

R. M. Sainsbury writes that quite a number of the many different views that have been called projectivist involve ‘the thought that some concepts are *defective* because they are “projections” of subjective human feelings’.³⁰⁶ Mark Johnston characterizes the projectivist view as holding that ‘value is not a genuine feature of persons, acts, states of affairs, etc., but only appears so because we *mistake* features of our evaluative responses for features of such things’.³⁰⁷ Talk of value concepts being defective, of value not being a genuine feature, of things appearing differently from how they really are because of some mistake, leads us to ask whether projectivism implies that our first order, evaluative thought, discourse and practices are globally and systematically in error. As we will see, some projectivists—following John L. Mackie—have argued that they are, whilst other projectivists such as Simon Blackburn have argued the opposite. Blackburn characterizes his quasi-realist programme as the enterprise of showing that ‘even on anti-realist grounds there is nothing improper, nothing “diseased” in projected predicates’.³⁰⁸

In the present chapter I discuss how value projection should best be understood. In the next and final chapter I proceed to discuss the implications of projectivism for theories of evaluative language. I there identify some of the alternative theories of evaluative language that may be

³⁰⁵ Rønnow-Rasmussen (2003) (p. 253). By contrast, Rabinowicz, in responding to an objection against the object interpretation, writes that ‘the object interpretation *presupposes* [my emphasis] a subjectivist (or “projectivist”) theory of value’ Rabinowicz & Österberg (1996) (p. 19).

³⁰⁶ Sainsbury (1998) (Section 1)

³⁰⁷ Johnston (1989) (p. 295)

³⁰⁸ Blackburn (1984) (p. 171)

accepted in conjunction with a subjectivist/projectivist theory of the nature of value.

II. Projection

We first need to isolate the kind of projection that would have philosophical relevance. In psychology, for example, projection is a form of defence mechanism in which a person displaces unwanted feelings onto another person.³⁰⁹ This is surely not the kind of projection that the philosophical projectivist has in mind. Within philosophy, many discussions of projectivism begin with David Hume's remark that

'Tis a common observation, that the mind has a great propensity to spread itself on external objects, and to conjoin with them any internal impressions, which they occasion, and which always make their appearance at the same time that these objects discover themselves to the senses.³¹⁰

Hume exemplifies this propensity by saying that 'as certain sounds and smells are always found to attend certain visible objects, we naturally imagine a conjunction, even in place, betwixt the objects and the qualities.' But though it comes naturally to human minds, this is nonetheless a mistake, since the qualities in question are 'of such a nature as to admit of no such conjunction, and really exist no where.' According to Hume, this same underlying psychological propensity explains 'why we suppose necessity and power to lie in the objects we consider, not in our mind, that considers them'. '[N]otwithstanding', he continues, 'it is not possible for us to form the most distant idea of that quality, when it is not taken for the determination of the mind.'³¹¹ Leaving aside for the moment whether this

³⁰⁹ Encyclopedia Britannica: <http://www.britannica.com/EBchecked/topic/478472/projection>

³¹⁰ *Treatise* (i.iii.xiv.)

³¹¹ *Ibid*

story is accurate for sounds, smells, colours and causal connections, we can ask whether it is accurate for values.

Objectivists have indeed been accused of having to postulate a ‘conjunction, even in place’ between objects and values; namely, a metaphysical supervenience relation holding between the latter and the non-evaluative features of the former. This is, for example, the underlying thought of one strand in Mackie’s argument from queerness:

Another way to bring out this queerness is to ask, about anything that is supposed to have some objective moral quality, how it is linked with its natural features. What is the connection between the natural fact that an action is a piece of deliberate cruelty—say, causing pain just for fun—and the moral fact that it is wrong? It cannot be entailment, a logical or semantic necessity. Yet it is not merely that the two occur together. The wrongness must somehow be ‘consequential’ or ‘supervenient’; it is wrong because it is a piece of deliberate cruelty. But just what *in the world* is signified by this ‘because’? And how do we know the relation that it signifies, if it is something more than such actions being socially condemned, and condemned by us too, perhaps through our having absorbed attitudes from our social environment? It is not even sufficient to postulate a faculty which ‘sees’ the wrongness: something must be postulated which can see at once the natural features that constitute the cruelty, and the wrongness, and the mysterious consequential link between the two. Alternatively, the intuition required might be the perception that wrongness is a higher order property belonging to certain natural properties; but what is this belonging of properties to other properties, and how can we discern it? *How much simpler and more comprehensible the situation would be if we could replace the moral quality with some sort of subjective response which could be causally related to the detection of the natural features on which the supposed quality is said to be consequential.*³¹²

Echoing Hume’s words, the value subjectivist could well insist that it is indeed hard for us to understand what values are supposed to be (and how they relate to the non-evaluative properties of the objects we value) unless we take it for the ‘determination of the mind’. On the subjectivist view, not

³¹² Mackie (1977) (p. 41). My emphasis.

just value, but also the supervenience of value on the natural features of objects is a product of our attitudes.

Simon Blackburn is another subjectivist who stresses the need for making this relation intelligible. And he thinks that projectivism does this:

[One] argument for a projective moral theory is in effect a development of the simple thought that moral properties must be given an intelligible connection to the natural ones upon which they somehow depend. It generates a metaphysical motive for projectivism.³¹³

This, in turn, is tied up with an appeal to ontological economy:

The projective theory intends to ask no more from the world than what we know is there—the ordinary features of things on the basis of which we make decisions about them, like them or dislike them, fear them and avoid them, desire them and seek them out. It asks no more than this: a natural world, and patterns of reaction to it.³¹⁴

On the subjectivist view, as we saw in Chapter 2, the values that supervene on ordinary features of objects are taken to be conferred by our attitudes, our ‘patterns of reaction’, towards the objects in question. If the relevant attitude is a preference, for example, we can say that value supervenes on precisely those non-evaluative properties of the object in virtue of which it is being preferred. This might be seen as taking at least some of the mystery out of supervenience by identifying an ‘intelligible connection’ between values and the non-evaluative properties upon which they depend. But insofar as the subjectivist can make this connection intelligible, this is done by introducing a further relation into the picture: namely, the grounding or constituting relation. The question remains whether the notion of projection adds anything to our understanding of this latter relation.

Value-Projection Understood in Analogy with Literal Projection

³¹³ Blackburn (1984) (p. 187)

³¹⁴ *Ibid* (p. 182)

One way to understand value projection would be in analogy with the kind of projection that is produced by a film projector or a spotlight. Such devices literally project things (light waves) onto other objects. But does this literal sense of projection shed any light (pardon the pun) on the idea of a constitutive ground of value? Richard Joyce doubts that anyone seriously believes that when a person projects an attitude, this attitude ‘literally flies forth from her brain and laminates the world’. In Joyce’s slogan, ‘Projecting attitudes is not like projectile vomit’.³¹⁵ Johan Brännmark (also discussing the analogy with flashlights and film projectors) similarly insists that it is quite obvious that we do not (using Hume’s words) ‘gild or stain’ the world with values in that way: ‘There are no magical value-beams emanating from our eyes’.³¹⁶ And such ‘value-beams’ would have to be rather ‘magical’ indeed, especially if value is thought to be projected upon abstract objects as well as concrete objects, and upon objects that are located in other possible worlds and at other times! No film projector or spotlight we know of is capable of such feats.

Brännmark also draws our attention to a related concern:

Clearly, if I ‘project’ values onto the world they are in some sense ‘there’ for me, but if I am not literally adding something to the world, then in what sense can such values be said to be ‘there’ for others as well?³¹⁷

If the values that are being projected onto an object by the attitudes of subject S_1 are supposed to be ‘there’, not only from the perspective of S_1 , but also for S_2, S_3, \dots, S_n , this would seem to require that S_1 *literally adds something to the world* that would then be there for others to see. Indeed, it seems to require that the attitude, in Joyce’s words, ‘literally flies forth from her brain and laminates the world’. How else could the values ‘be there’ for others as well?³¹⁸ If S_1 and S_2 share a common attitude, the value thus constituted is, of course, there for both of them, but this does not require

³¹⁵ Joyce (2009) (p. 60)

³¹⁶ Brännmark (2000) (p. 66)

³¹⁷ Ibid (p. 69)

³¹⁸ Recall also the objections I made against the ‘valuable until further notice’ view at the end of in Chapter 3. This view invites a literal understanding of projection.

the product of S_1 's projection to be there (as in any way distinct from the product of S_2 's own projection) to be seen by S_2 . Each individual subject views the valuable object from his or her own perspective, through the lens of his or her own attitudes. Similarly in cases where S_2 identifies with S_1 and takes over the latter's attitudes and makes them his own, as recommended by the participant model. When S_2 has taken over S_1 's attitudes, these attitudes are then treated as S_2 's own in the latter's evaluative thinking. There thus seems to be no need for us to assume that S_2 would be able to perceive S_1 's projections in any other, more direct, sense. Furthermore, if we would nonetheless assume that the values projected onto the world by one subject could somehow be perceived by another, we would then need to explain also how the products of the former's attitudes could possibly cause something in the latter. Such a project does not seem promising.

Value requires the first-person perspective. Subjectivists hold that values can only be seen through the lens of the subject's pro and con attitudes (some of which might perhaps have been taken over from other subjects in accordance with the participant model). What can be seen from the third-person viewpoint, I have argued, are only different subjects' diverse and often conflicting attitudes and the respective objects of their attitudes as well as the non-evaluative features of the objects in virtue of which they are being favoured or disfavoured. We can infer from this information that value-constitution is going on, and we can derive non-evaluative conclusions of the type 'Object O is good relative to subject S_1 ', but we cannot reach any evaluative conclusion such as 'O is good' or 'O is good for P'. Now, given this picture, assuming something further—something more directly analogous to the images produced by film projectors—seems to me to add nothing useful to the theory and indeed would only invite further troublesome questions.

Yet some projectivists seem to be committed to a rather literal understanding of value projection. R. M. Sainsbury has, for example, argued for an interpretation of Hume that takes Hume's idea that taste has a 'productive faculty' and 'raises in a manner a new creation' quite literally.³¹⁹ We would then, it seems, have to think of the products of value projection as being as real as the light waves produced by a spotlight (in sharp contrast

³¹⁹ Sainsbury (1998)

to causal connections that Hume thought ‘really exist no where’³²⁰). In following Sainsbury’s Hume, we would presumably satisfy the aforementioned demand put forward by Brännmark. As mentioned in Chapter 2, Brännmark asks in what sense we are justified in speaking of value as in any way grounded on the subjectivist view. He specifies what he demands by saying that ‘the question is whether there really is a metaphysical substance “value” which has some kind of “source” [the constitutive ground] from which it can be “conferred” [upon objects], or instead whether ‘it is just a metaphysical-sounding way of expressing something that is not metaphysical, but rather a psychological phenomenon’.³²¹ But, as I insisted, though it is perfectly in order to require that the subjectivist tells us more about the relation between valuable objects and the attitudes on which their values are thought to be grounded, demanding that this be done in a way that treats value as a ‘metaphysical substance’ seems to beg important questions. I think that the subjectivist should be happy to understand value as a psychological phenomenon and eschew any idea of value being a metaphysical substance. The idea that value is a product of attitudes should not be understood literally.

But even if one were to insist that value be treated as a metaphysical substance that is literally projected by the subject upon the value bearer, this assumption does not seem to do any explanatory work. Rather it introduces further things that need explaining. It also seems to undermine what Blackburn referred to as a ‘metaphysical motive’ for projectivism: the potential to make the relation of supervenience intelligible. It is hard to

³²⁰ Sainsbury (ibid) writes ‘Projection in the case of necessity is errorinfected; in the case of morals our projections (if they deserve to be so called) are faultless. We must make room for two kinds of mental operation: when Hume talks, in connection with necessity, of the mind’s “great propensity to spread itself on external objects”, he is talking about an operation which leads to error; but if we allow that his view of morals is adequately expressed in terms of the mind’s “productive faculty . . . [its] gilding and staining all natural objects with the colours, borrowed from internal sentiment”, then we must think of gilding and staining as quite different from spreading, since these processes do not of themselves lead to error’ (Section 1). On the face of it, Sainsbury’s interpretation seems to clash with Hume’s insistence (previously quoted) that when you examine an action allowed to be vicious, you don’t find that ‘matter of fact, or real existence, which you call vice’, and that ‘in which-ever way you take it, you find only certain passions, motives, volitions, and thoughts’.

³²¹ Brännmark (2002) (p. 67)

understand how projected values, understood in a literal way, attach themselves to the ordinary features of the things we value.

Summing up, we now have several reasons for doubting that the analogy with film projectors and spotlights, taken literally, will help the subjectivist to illuminate the idea of a constitutive ground of value. Given that values can be projected upon abstract objects such as events and states of affairs, as well as objects located in other possible worlds and at other times, and that the products of value projection are 'there' only from the first-person perspectives of the individual subjects in question who possess the relevant conative attitudes, the analogy of projection should not be taken literally.

Value-Projection Understood in Analogy with Secondary Qualities

Leaving aside the literal interpretation of projection, we might look for other analogies. One candidate is an analogy between value and so-called secondary qualities like colour. As pointed out by Stephen Darwall, the analogy between value and colour, especially a projectivist approach to both, goes back to Thomas Hobbes. Darwall argues for an interpretation of *Leviathan* as exhibiting an application to the case of value of a projectivist approach to colour that Hobbes took from Galileo.³²² Darwall further says that *Leviathan* provides one of the first, if not *the* first, expressions of a projectivist meta-ethics, for reasons similar to those that underlie projectivism today.³²³

According to Darwall, Hobbes thought that

the proper place to locate ethical thought is in practical reasoning, from the agent's point of view, but also that ethical or normative thought and discourse are expressions or projections of what, in Hobbes's view, makes us deliberating agents in the first place, our *desires*. In desiring [some object] . . .

³²² Darwall (2000) (p. 321)

³²³ *Ibid* (p. 319)

we ascribe to it a property, that of being good (something we ought to achieve), that it does not literally have.³²⁴

The agent has these normative thoughts because she has desires. They are the ‘appearances’ of her desires. [But] there is nothing in the nature of the objects of her desires that answers to the normative properties she attributes to them in having desires.³²⁵

So far so good, but what has it got to do with colour? The analogy is this:

In color experience we see objects as colored, as though they had an objective, categorical color property. But they actually have no such property.³²⁶

The analogy between colour and value can also be found in Hume’s writings. After having said (1) that vice and virtue are not to be found in the objects, but instead in ourselves³²⁷, and that (2) when you call something vicious or virtuous you mean nothing but that you have a sentiment of blame or approbation toward it, Hume immediately goes on to say (3) ‘Vice and virtue, therefore, may be compar’d to sounds, colours, heat and cold, which, according to modern philosophy, are not qualities in objects, but

³²⁴ Ibid (p. 318)

³²⁵ Ibid (p. 333)

³²⁶ Ibid (p. 320)

³²⁷ Hume wrote in the *Treatise*: ‘The vice entirely escapes you, as long as you consider the object. You never can find it, till you turn your reflexion into your own breast, and find a sentiment of disapprobation, which arises in you, toward this action. Here is a matter of fact; but ’tis the object of feeling, not of reason. It lies in yourself, not in the object’ (iii.i.i.). On a subjectivist reading of this claim, what lies ‘in yourself’ is not the vice as such, but only the ‘sentiment of disapprobation’ that is identified with the constitutive ground of the vice. The supervenience base of the vice, the non-evaluative properties towards which the sentiment arises, is accurately located ‘in the object’. The vice itself, however, Hume held, is not a ‘matter of fact, or real existence’, and this is the reason why you cannot find it as long as you consider the object from an objective, third-person point of view. It is only when you ‘turn your reflexion into your own breast’, and thus come to view the situation from your own first-person perspective that you will be able to see the vice. It can only be seen by those who share the relevant attitudes upon which it is grounded.

perceptions in the mind'.³²⁸ At least on one interpretation, the first of these three claims expresses a subjectivist view of value of just the type that I characterized in this work. As such, it is a claim about the nature of value. The second, on the other hand, is a semantic claim, a claim about the nature of evaluative language, the function of value judgements and the meaning of evaluative words and expressions. My subjectivist is not committed to this second claim or to any particular view about evaluative language. Such views will be discussed in the next chapter. But what kind of claim is the third one?

I think that though there undoubtedly is something to this analogy it also has a potential to mislead. It risks making value too much a matter of perception. The same analogy is also heavily used by defenders of so-called dispositional theories of value, and Marc Johnston points out that the analogy with secondary qualities, and with colour in particular, has been used to undermine rather than vindicate projectivism about value.³²⁹ Blackburn points to a number of disanalogies between value and secondary properties that to him seem great enough for 'it to appear a severe error of philosophical taste' to expect a theory of evaluation to look very much like a theory of secondary quality perception.³³⁰ However, this analogy would deserve a much more thorough discussion than I have given it here.

Value-Projection Understood in Analogy with the Pathetic Fallacy

Mackie, in introducing his discussion of projection in *Ethics*, appeals to an analogy with the 'pathetic fallacy': 'the tendency to read our feelings into their objects'. The example he gives is of a fungus that fills us with disgust and thereby inclines us to 'ascribe to the fungus itself a non-natural quality of foulness'.³³¹ But, as Mackie goes on to argue, there is much more to value projection than this. Richard Joyce comments that Mackie's is actually not a

³²⁸ Ibid

³²⁹ Johnston (1989) (p. 295)

³³⁰ Blackburn (1985) (pp. 13–15)

³³¹ Mackie (1977) (p. 42). The same example is used by Mackie in his (1946).

good example of the pathetic fallacy. Committing the pathetic fallacy proper would involve attributing to the fungus, not the quality of foulness, but the human trait of being disgusted. What Mackie refers to by the disjunctive phrase ‘projection or objectification’ is not a matter of ‘reading our feelings into their objects’ at all, but rather a process by which, in Joyce’s words, ‘our feelings cause us to read into their objects qualities that we would not otherwise judge them to have’.³³² Joyce also notes that Mackie, in his 1980 book *Hume’s Moral Theory*, more accurately writes that we ‘read some sort of image of [our] sentiments’ into the objects that arouse them, leading us to think that these objects possess qualities that they don’t really possess. Also, in a 1946 paper, Mackie wrote that ‘in objectifying our feelings we are also turning them inside out . . . The feeling and the supposed quality are related as a seal or stamp and its impression’.³³³ The foulness of the fungus would thus be ‘some sort of image’ of our disgust, and, as Joyce suggests, moral wrongness is ‘some sort of image’ of our disapproval.³³⁴ In line with this, a value subjectivist would hold that a positive value such as goodness is some sort of image of a pro attitude, and negative values some sort of image of our con attitudes. This analogy is surely better than the preceding ones. Can we do still better?

The Chess Analogy

In order to illustrate the difference between the supervenience base and the constitutive ground of value, Rabinowicz introduces an analogy with the conventions of chess: ‘[T]hat a particular move in chess is admissible supervenes on the internal properties of the move together with the internal properties of the game situation. But the constitutive ground of its admissibility is to be found in the social convention that lies behind the rules of chess’.³³⁵ In a separate paper from the same year, the same author notes that the chess analogy is not fully satisfactory. The dependence of

³³² Joyce (2010) (pp. 2–3)

³³³ Mackie (1946) (pp. 81–82) as quoted by Joyce (2010).

³³⁴ Joyce (2010) (p. 3)

³³⁵ Rabinowicz & Österberg (1996) (p. 21)

value on its constitutive ground differs from the corresponding dependence relation in the chess case. If the chess convention had been different, a particular move that is now permissible might not have been so. But if our preferences had been different, values would nonetheless remain the same (as was argued in Chapter 3 above). This disanalogy has to do with values' unlimited range of application or modal scope. While values are products of actual attitudes, they are applicable even to possible worlds in which our attitudes are different.³³⁶

Projection as Metaphor

In my discussion so far we seem to have a proliferation of metaphors: the mind's propensity to spread itself on external objects, gilding, staining, conferring, bestowing, reading some sort of images of feelings into objects, raising in a manner new creations, etc. Joyce points out that discussions of projectivism have rarely given proper attention to the fact that the theory is generally offered and thought about in metaphorical terms. He also asks how we can determine whether a metaphor is adequate, but he provides no answer to this question.³³⁷

We might agree that projection is best seen as a metaphor, but one potential risk with treating it merely as a metaphor is that we might perhaps end up with a merely metaphorical explanation (of perhaps merely metaphorical values)!

Projection as Empirical Hypothesis?

Joyce is not satisfied with metaphors. He writes 'Evidently, projectivism is a theory in need of translation into literal terms before it can be properly

³³⁶ Rabinowicz (1996) (p. 32). This insight is attributed to Sven Danielson. Yet, despite the recognition of its not being perfect, the chess analogy resurfaces in Rabinowicz & Rønnow-Rasmussen (2000) (p. 37). The same analogy is also discussed in Rønnow-Rasmussen (2003) and Rønnow-Rasmussen (2011).

³³⁷ Joyce (2009) (p. 60)

assessed.³³⁸ Joyce then goes on to provide such a translation. He begins by characterizing what he calls minimal projectivism as a conjunction of two subtheses: one phenomenological and one about the causal origin of this phenomenology:

- (1) 'We experience [value] as an objective feature of the world'³³⁹
- (2) 'This experience has its origin in some non-perceptual faculty; in particular, upon observing certain actions and characters (etc.) we have an affective attitude (e.g., the emotion of disapproval) that brings about the experience described in 1'³⁴⁰

³³⁸ Ibid

³³⁹ Ibid (p. 56). To better suit my purposes, I have replaced 'moral wrongness (e.g.)' with 'value'. It should be noted (and as Joyce is well aware) that not all values are experienced and that we can, and often do, reach evaluative judgments by other means than immediate observation of the object under evaluation, for example by reasoning. However, it is plausibly held that the most basic evaluative judgments come about by direct experience and that in the absence of any such direct experiences of value, there would be nothing to reason about in evaluative matters. Moreover, since Joyce insists that 1 is an empirical claim that can be tested as such, one would expect a formulation that eschews technical philosophical terminology. (Joyce is, of course, not blind to this problem and has a whole section of his paper where he discusses how the metaphor of value projection can be turned into a testable empirical hypothesis.) In order to test empirically if 1 really is an apt description of our phenomenology, Joyce would first have to inform his test subjects about the subjective/objective distinction, which seems to require them to become amateur philosophers. Further, even assuming that we have the objective/subjective distinction laid out clearly before us, we might still rightly wonder what it is to experience something *as objective*. Furthermore, though many projectivists would probably be attracted to a claim like 1 one could deny that some such claim forms a necessary component of a projectivist view in the first place. 1 is, after all, a phenomenological thesis and projectivism need not be seen as including such a thesis.

³⁴⁰ Ibid (p. 56)

Joyce insists that both 1 and 2 are to be regarded as empirical theses that (at least in principle) can be tested as such.

III. Projectivism and the Ontological Status of Value

Joyce's minimal version of projectivism does not by itself commit one to any particular view about the ontological status of values (realist or anti-realist) and is thus compatible with the literal interpretation of projection discussed above and attributed to Hume by Sainsbury. Joyce explicitly characterises the view of Sainsbury's Hume as a realist form of projectivism. However, it is not clear that acknowledging literal projection is sufficient to be a value realist. This depends on how we more precisely define realism. Understood in this minimal way, projectivism is thus not anti-realist per definition. There seems to be conceptual room for a realist projectivism. Presumably, on such a view, values would be mind-dependent only in the sense that Human-made artefacts like tables and chairs are. Values would be products of attitudes, but once created they would not need to be sustained continuously by attitudes, but would gain some kind of attitude-independent existence. Such a view, I believe, would have many problems and few advantages. Most projectivists from Hobbes and Hume onwards through Mackie, Blackburn and Joyce, have clearly tended towards anti-realism and projectivism's ready compatibility with a rejection of realism should be considered one of its advantages. Philosophers who are sceptical about value realism should therefore find projectivism attractive.

One might insist that realism versus anti-realism constitutes a false dichotomy. Russ Shafer-Landau distinguishes three different types of view about the reality of a given domain: realist, anti-realist (or nihilist) and constructivist. Nihilists reject the reality of the domain, while both realists and constructivists endorse it. However, in opposition to realists, constructivists

explain this by invoking a *constructive function* out of which the reality is created. This function has moral reality as its output. What distinguishes constructivist theories from one another are the different views about the proper input. Subjectivists [who are one among several types of constructivists mentioned by Shafer-Landau] claim that individual tastes and opinions are the things out of which moral reality is constructed. . . . what is common to all constructivists is the idea that moral reality is constituted by the attitudes, actions, responses, or outlooks of persons, possibly under idealized conditions.³⁴¹

If we replace ‘moral reality’ with ‘evaluative reality’ or simply ‘value’ (and perhaps ‘tastes and opinions’ with ‘pro and con attitudes’), Shafer-Landau’s characterization of constructivism seems to pretty closely capture the value subjectivism that I have given in previous chapters. It would seem then that, if we allow constructivism as a third option in addition to realism and anti-realism/nihilism, value subjectivists are or at least can be constructivists in the present sense of the term. One could perhaps reason that if taste has a productive faculty, as Hume put it, capable of raising a (‘in a manner’) new creation, or if, as Gauthier put it, values are products of our affections, then these new creations or products must be allowed some kind of reality. Perhaps then a projectivist understanding of the subjectivist’s notion of value-constitution can be thought of as being compatible with endorsing the reality of value in a constructivist sense, or at least that constructivism remains a possible alternative for the projectivist/subjectivist.

The constructivist alternative will probably be more attractive to subjectivists who are universalist absolutists than to individualist relativist subjectivists. This has been noted by Johan Brännmark who suggestively asks: ‘[M]ight one not conceive of a social projectivism that can draw on some of the advantages of understanding the constitution of values as being done by a plural, rather than an individual, subject?’³⁴² ‘Social constructivism’, Brännmark holds, ‘provides us with a sense in which we do actually literally add something to this world’.³⁴³ Brännmark further thinks

³⁴¹ Shafer-Landau (2003) (p. 14). Emphasis in original.

³⁴² Brännmark (2002) (p. 95, n. 52)

³⁴³ Ibid (p. 95)

that, unlike other forms of projectivism, social constructivism manages to make good on its metaphor. I, however, doubt that this is so, especially in light of the absence of an account of what exactly it is that we are thought to be literally adding to the world and the process by which this addition is supposed to work.

Joyce goes on to characterize what he calls metaphysical projectivism, which is a form of projectivism that goes beyond minimal projectivism by adding a third claim to 1 and 2 above:

(3) 'In fact, [value] does not exist in the world'³⁴⁴

This third claim entails anti-realism. In 1 Joyce specified that we experience value as an objective feature of the world, so perhaps we should modify 3 to include this qualification as well (in case it is thought to make a difference):

(3') In fact, value is not an objective feature of the world

A potential worry with metaphysical projectivism/anti-realist subjectivism could be that there is a tension, or perhaps a downright contradiction, between, on the one hand, the main subjectivist claim that values are constituted by (are the products of) attitudes, and on the other hand the anti-realist claim that values do not exist. What sense can be made of the idea of values as products in the light of the idea that values are not real? Are values to be thought of in such oxymoronic terms as non-existent products? This seems scarcely intelligible.

I, however, believe that this supposed tension is merely apparent. To be good relative to a subject S is not to possess the property of goodness, but rather to have a certain status from S's point of view as determined by S's actual, present conative attitudes. I have argued that by favouring an object O in virtue of some of its ordinary, non-evaluative features, S confers upon these features the status of being good-makers relative to S. There is really no need to understand this conferring or bestowing of a status as involving any kind of metaphysical substance being literally brought into existence by

³⁴⁴ Joyce (2009) (p. 56)

the subject and somehow transferred to the object of her attitude. It might well be that the subject would come to see the object *as if* it had this status *objectively*—as if it was good not just relative to her point of view but from the third-person perspective. But since subjectivists claim that this would involve a mistake, there does not seem to be any need on this account for such value theorists to allow that values exist (neither in the realist sense or in the constructivist sense). Indeed, as I've insisted above, I believe that a subjectivist who would maintain that projection should be understood literally has a lot more explaining to do. What sense could be made of the idea that by taking an attitude towards something you thereby literally bring something into existence that somehow attaches to the object of your attitude? I doubt that such problems could be overcome.

Without claiming to have ruled out either realism or constructivism once and for all, I think that it is fair to say that, in the light of the above considerations, the burden of proof lies with those who would want to combine subjectivism with a realist or constructivist view of the ontological status of value. Subjectivism seems naturally to invite the third subthesis (or something like it) of Joyce's metaphysical projectivism.

Does Joyce's metaphysical projectivism add anything of substance to the subjectivist view outlined in the previous chapters? Yes and no. Something along the lines of the genealogical thesis expressed in 2 will undoubtedly form an important premise in an argument for subjectivism. If it can be explained how value appears to be an objective feature of the world, without assuming that it is an objective feature of the world, then we seem to have undermined an important presumption in favour of objectivism.³⁴⁵ Insofar as the objective appearance of value is taken as (*prima facie*) evidence for objectivism, this evidence is undermined by a successful explanation of this type. If objective values don't play any causal role in the explanation of our attitudes and experiences, then we seem to lack evidence for objective values.³⁴⁶ However, fleshing out this argument would require a book of its own and is not my present concern. What I am exploring here is whether projectivism can elucidate the idea of a constitutive ground of value, and Joyce's projectivism seems to add little to that project.

³⁴⁵ Don Loeb (2010) has argued that there no such presumption in the first place.

³⁴⁶ An argument of this kind can, for example, be found in Gilbert Harman (1977) and expanded upon by Joyce (2006).

This concludes the main part of this work. The next chapter will explore implications of accepting value subjectivism as outlined above with respect to evaluative language. Relating value subjectivism to theories of evaluative language and thus putting subjectivism in a wider meta-ethical perspective might however shed some light back onto subjectivism itself.

8. SUBJECTIVISM AND EVALUATIVE LANGUAGE: THE ERRORS OF OUR WAYS

Moore is right, I think, in holding that when we say a thing is good we do not *merely* mean that we have towards it a certain feeling, of liking or approval or what not. There seems to me no doubt that our ethical judgments claim objectivity; but this claim, to my mind, makes them all false.

Bertrand Russell, *Is There an Absolute Good?*

I. Introduction

Assuming we accept the subjectivist view about the nature of value outlined in the foregoing chapters is, what implications does this have with respect to theories about evaluative language? Are there any restrictions concerning which views of the latter kind that can be held in conjunction with a subjectivist view about the nature of value? Given the idea that value is grounded on attitudes, is there some theory of the semantics of evaluative language that is more fitting than other such theories, or is the semantic field wide open for value subjectivists? These questions will be explored in the present chapter. Three families of semantic theory will be considered: naturalistic theories, non-cognitivist theories and error theories. It should be noted at the outset that the discussion that follows will not do justice to the rich variation of positions within these three theoretical families but instead

will focus on broad features. Hopefully my discussion will still give some general idea as to which considerations that are raised by value subjectivism in relation to the nature of evaluative language; and, by extension, to evaluative truth and knowledge.

II. Subjectivism and Naturalism

As noted in Chapter 1, the term *subjectivism* is often used to refer to the view that evaluative utterances are used to report the speaker's own attitudes. When Stephen Darwall argues that, contrary to common interpretation, Thomas Hobbes was not a subjectivist but instead a projectivist, he is using subjectivism to refer to this view of what evaluative statements mean.³⁴⁷ The way I have been using subjectivism in this work is much closer, perhaps identical, to how Darwall uses projectivism. Thus, if Darwall's interpretation of Hobbes is correct, Hobbes was a subjectivist in my sense of the term. It should be clear that the value subjectivist is not committed to this view of evaluative language, but can it even be held in conjunction with it? One could think that, since subjectivism (the semantic theory) and projectivism are about different things, they must be compatible and hence that Hobbes could have been both at once. But this is not entirely obvious. The view that evaluative utterances report the speaker's attitude seems to make value facts into facts about psychology—ordinary, natural facts about a speaker having a certain attitude (or being disposed to have this attitude, or some such thing)—and it is not easy to see what role is left for anything of a projective nature to play in this story.

David Lewis characterizes his own dispositional theory of value as 'subjective' in that it 'analyses value in terms of attitudes'.³⁴⁸ On Lewis's view, 'values are what we are disposed to value'.³⁴⁹ More precisely, Lewis

³⁴⁷ Darwall (2000) (pp. 315–317)

³⁴⁸ Lewis (1989) (p. 113)

³⁴⁹ Ibid

analyses value in terms of second order desires; as what we ‘desire to desire’. However, as I will argue presently, his view is not subjectivist in my sense of the term. First, Lewis’s analysis is semantic, whilst (as argued in Chapter 1) my subjectivist offers no semantic analysis. In his own words, Lewis ‘advances an analytic definition of value’ on which ‘to be a value—to be good, near enough—means to be that which we are disposed, under ideal conditions, to desire to desire’.³⁵⁰ Secondly, and more to the present point, whilst Lewis’s theory ‘reduces facts about value to facts about our psychology’³⁵¹, my subjectivist refuses any such reduction. The realm of value is not a realm of objective fact, psychological or otherwise. There are, however, on the subjectivist picture, facts about what is valuable relative to this or that perspective. But such facts, I have insisted, are no more evaluative or normative than are facts about what is valuable according to someone. Thirdly, because of its reduction of evaluative claims to psychological claims, Lewis’s theory straightforwardly ‘allows us to seek and gain knowledge about what is valuable’.³⁵² On subjectivist views, as I will argue, value is not a matter of knowledge, at least not the ordinary kind of knowledge that we have about other things. Knowing that something is such that you desire to desire it (under some special conditions), is, *pace* Lewis, not the same as knowing that it is good. Even though it analyses value in terms of attitudes, Lewis’s theory does not claim that values are constituted by attitudes, hence is not a form of subjectivism in my sense.³⁵³

Another candidate naturalistic definition of *good* would be to let it mean the same as ‘good relative to’, which I argued in Chapter 5 should be understood as good from the perspective of subjects and with the relevant perspective fixed with reference to the actual, present conative attitudes of the individual subject). This definition of *good* would give it an explicitly indexical character, incorporating into the meaning of the word an essential reference to the speaker’s attitudinal perspective. This proposal would also,

³⁵⁰ Ibid (p. 116)

³⁵¹ Ibid (p. 113)

³⁵² Ibid

³⁵³ Other dispositional accounts of value have been offered by Marc Johnston (1989) and by Michael Smith (1989). On such views, value, or goodness, is a natural, dispositional property, and as such straightforwardly objective.

like that of Lewis, constitute a reduction of an evaluative term to a non-evaluative, natural term. And, again like that of Lewis, this would give us evaluative truth and knowledge.³⁵⁴

Mackie pointed out that naturalism is in accordance with the view that evaluation is more a matter of knowledge and less a matter of choice than any non-cognitive analysis allows. However, in satisfying this demand, Mackie claims, naturalism introduces a converse deficiency: 'On a naturalist analysis, moral judgments can be practical, but their practicality is wholly relative to desires or possible satisfactions of the person or persons whose actions are to be guided; but moral judgments seem to say more than this.'³⁵⁵

Lewis adds that his view is naturalistic also in the sense that 'it fits into a naturalistic metaphysics. It invokes only such entities and distinctions as we need to believe in anyway, and needs nothing extra before it can deliver the values'. But it 'delivers' the values only by identifying them with natural properties, by reducing value facts to ordinary, natural facts. Is such a reduction plausible? Lewis himself seems to doubt its adequacy, admitting that in 'requiring values to be all that we might wish them to be, we bring on the error theory.'³⁵⁶

Strictly speaking, Mackie is right: genuine values would have to meet an impossible condition, so it is an error to think there are any. Loosely speaking, the name may go to a claimant that deserves it imperfectly.³⁵⁷

³⁵⁴ As noted, Lewis's theory allows us to seek and gain knowledge about what is valuable. 'This knowledge', Lewis continues, 'is *a posteriori* knowledge of contingent matters of fact. It could in principle be gained by psychological experimentation. But it is more likely to be gained by difficult exercises of imagination, carried out perhaps in a philosopher's or novelist's armchair' (p. 113). If we reduced 'good' to 'good relative to', we would similarly have been able to seek and gain knowledge about what is valuable by empirically investigating the attitudes of subjects. I don't find that very plausible for reasons I will elaborate on shortly.

³⁵⁵ Mackie (1977) (p. 33). Mackie's discussion focuses on ethics and specifically moral requirements but I believe his reasoning is applicable to evaluative language more generally.

³⁵⁶ Lewis (1989) (p. 134)

³⁵⁷ Ibid (pp. 136–137)

The question is why we would want to speak loosely when we can speak strictly. Lewis thinks that ultimately where we stand on this issue is a matter of temperament, and the error-theorist Joyce seems to agree with him³⁵⁸.

Had my ambition been to establish that we should reject semantically naturalistic views, then I would at this point have had to pay attention to the variation in naturalistic views and discuss a number of familiar arguments and counterarguments. But my aim here is much more local. Whatever we might think of naturalistic analyses of the semantics of evaluative language *on their own merits*, we may ask whether they are compatible with value subjectivism as I have characterized it. If goodness is an ordinary, natural property—although perhaps a dispositional or otherwise complex property—it is difficult to see how it can at the same time be thought of as being a projection of attitudes. Furthermore, if value is natural it would not seem to require the first-person perspective in order to be seen, which I have argued is an essential claim of value subjectivism. To say that value subjectivism and semantically naturalistic views cannot be combined—that they are strictly speaking incompatible—may be a stronger conclusion than this quick and rough argument warrants. I therefore rest content with the more modest conclusion that there seems to be at least a certain tension between the two views.

III. Subjectivism and Non-Cognitivism

As previously mentioned, John Mackie and Simon Blackburn are both projectivists, but they hold opposing views about the semantics of evaluative language. I will try to tease apart semantic debates from the debate concerning projectivism. But this is not an entirely easy task however, since both Mackie and perhaps especially Blackburn see projectivism as intimately

³⁵⁸ See Joyce (2012)

connected with their respective general meta-ethical projects: Mackie with his error theory and Blackburn with his quasi-realism.³⁵⁹

Projection is a recurring theme in Blackburn's writings.³⁶⁰ He says that it is when 'we speak and think as though there were a property of things which our sayings describe, which we can reason about, know about, be wrong about, and so on'³⁶¹ that we project an attitude onto the world. Blackburn thus seems to be adhering to Hamlet's declaration that 'there is nothing either good or bad, but thinking makes it so', but on the subjectivist view as I have characterized it, it is not by speaking or thinking that we project value onto the world, but rather in having an attitude.³⁶² However, Blackburn also speaks of 'a subjective source for value in itself',³⁶³ which is already much more closely aligned with how I have described subjectivism.

In Mackie's writings the term *projection* occurs rarely, but when it does he uses it to mean the same thing as *objectification*, the latter being a term he uses extensively.³⁶⁴

³⁵⁹ He does however add in a footnote 'It is important to be clear about the distinction between projectivism and quasi-realism. Projectivism . . . says that evaluative properties are projections of our own sentiments (emotions, reactions, attitudes, commendations). Quasi-realism is the enterprise of explaining why our discourse has the shape it does . . . if projectivism is true. It thus seeks to explain, and justify, the realistic-seeming nature of our talk of evaluations—the way we think we can be wrong about them, that there is a truth to be found, and so on. One might believe that quasi-realism is successful, yet still dislike projectivism, and one might like projectivism, but still believe that ordinary features of our thought are not explicable, quasi-realistically, but indeed involve error' (1984) (p. 180). The latter type of projectivism is, of course, that of Mackie and Joyce. Blackburn further thinks that quasi-realism 'removes the most important range of objections to projectivism—namely, that it cannot account for the phenomena of ordinary moral thinking' (ibid).

³⁶⁰ See, for example, Chapter 6 of his (1984). Blackburn too, like Mackie, invokes Hume's dictum of the mind's 'propensity to spread itself on the world and to 'gild or stain' natural objects with internal sentiments and he treats Mackie as a fellow projectivist in his (1984) and in his (1985).

³⁶¹ The book from which the quoted passage is taken is aptly titled *Spreading the Word*.

³⁶² Gauthier also refers to Hamlet's dictum and comments that 'thought it not the activity that summons value into being. . . . Desire, not thought, and volition, not cognition, are the springs of good and evil' (1986) (p. 21).

³⁶³ Blackburn (1984) (p. 198)

³⁶⁴ Mackie (1977) (p. 42)

On a subjectivist view, the supposedly objective values will be based in fact upon attitudes which the person has who takes himself to be recognizing and responding to those values. If we admit what Hume calls the mind's 'propensity to spread itself on external objects', we can understand the supposed objectivity of moral qualities as arising from what we can call the projection or objectification of moral attitudes.

In his later book *Hume's Moral Theory*, the disjunctive phrase 'projection or objectification' occurs again and what Mackie calls 'the objectification theory' is not only ascribed to David Hume but is also claimed to be 'very largely correct'.³⁶⁵ The objectification theory says that

[t]he meaning of moral statements is approximately [that an objective moral quality is found in a certain action (etc.); they are capable of being simply true or false], but the features ascribed to actions (etc.) in the distinctively moral (categorically imperative) part of these statements are fictitious, created in thought by the projection of moral sentiments onto the actions (etc.) which are the objects of those sentiments.³⁶⁶

Here it seems as if a certain view about the meaning of moral statements is taken by Mackie to be part of the objectification theory/projectivism. Hence, Mackie too, like Blackburn, seems to view projectivism as intertwined with issues of meaning in general and with his own meta-ethical position in particular. Since Mackie and Blackburn hold opposed views about the nature of evaluative language, they cannot both be right about these views being part of, or gain (unique) support from, the projectivist/subjectivist view of the nature of value that I take both of them to accept.

At one point Blackburn characterized his view in direct opposition to that of Mackie: 'Mackie believed that our ordinary use of moral predicates involved an error, because the underlying reality was as the expressive view [the type of view to which Blackburn adheres] claims, whilst in using those

³⁶⁵ Mackie (1980) (p. 72). For a detailed discussion of the role of projection/objectification in Mackie's overall argument, see Joyce (2010).

³⁶⁶ Mackie (1980) (p. 74)

concepts we claim more.³⁶⁷ Specifically, Mackie thinks that ordinary evaluative judgments involve a ‘claim to objectivity’, an assumption that there are objective values in just the sense in which he denies that there are.³⁶⁸ Blackburn agrees with Mackie that there are no objective values, but he rejects the further claim that an assumption to this effect is built into our everyday evaluative language. Hence, they agree about the subjective nature of value but disagree about the semantics of evaluative language.

On Mackie’s view, the claim to objectivity is ingrained in our evaluative language in the sense that it has been incorporated into the very meanings of evaluative (and normative) terms.³⁶⁹ Ordinary evaluators, in using the evaluative terms whose meaning in essence involves the mistaken assumption to the effect that values are objective, are thus globally and systematically in error in their evaluative discourse and judgments. Insofar as this is so, evaluative judgments essentially ascribe objective value to objects, and since there are no objective values to make such judgements true, no evaluative judgement comes out as true.

Blackburn is concerned to avoid Mackie’s conclusion. He sees his quasi-realism as the enterprise of showing, on anti-realist grounds, that there is no such global mistake in our evaluations as diagnosed by Mackie. Blackburn concedes, however, that it very well may be correct that ordinary evaluators ‘falsely imagine a kind of objectivity for values, obligations, and so on’, but he insists that it does not follow from this ‘that the error infects the practices of [evaluation], nor the concepts used in ways defined by that practice.’³⁷⁰ Hence, even though he denies that it effects our first order beliefs and practices, Blackburn too, like Mackie, diagnoses at least one error at the second order level: namely, the error of believing that there are objective values.

Sainsbury’s interpretation of Hume also stresses the difference between error at the second order level and error on the first order level:

³⁶⁷ Blackburn (1984) (p. 171)

³⁶⁸ Mackie (1977) (p. 35)

³⁶⁹ Ibid

³⁷⁰ Blackburn (1985) (p. 3). Blackburn writes ‘moralizing’, but I have replaced it with ‘evaluation’ to better suit my purposes.

Hume believed that our value concepts and discourse are perfectly in order; . . . there is no global first order error in these areas; . . . However, for Hume there are various errors at the reflective or philosophical level. Some people think [for example] of moral distinctions as made by reason rather than by feeling; . . . These errors do not, however, of themselves lead to error in our evaluative judgements.³⁷¹

Sainsbury can perhaps gain some support for his interpretation of Hume from the latter's insistence that, after having acknowledged that vice and virtue may be compared to sounds, colours, heat and cold, we observe that this 'has little or no influence on practice'.³⁷² This could be read as implying that though people are deeply and fundamentally mistaken about the nature and metaphysics of value, they are not necessarily similarly mistaken in their everyday evaluative thought, judgments and practices.

All value subjectivists agree that insofar as this error is widespread, it follows that many people are mistaken about the nature of value. But whilst for Mackie and other error theorists this error spreads and 'infects' (as Blackburn put it) our first order evaluative discourse, practices and beliefs, because the false view of the nature of value has become essentially built into the meanings of evaluative terms, Blackburn seems to believe that the error can be contained on the second order level and does not impact our substantial, first order evaluative discourse and practices.

Jan Österberg writes

The immanent [first-person] perspective only seemingly gives us access to objective values. Such a perspective gives rise to a systematic misapprehension: it makes us believe that the values we acknowledge in that perspective exist independently of us, whereas in fact they are projections of our preferences.³⁷³

³⁷¹ Sainsbury (1998) (Section 1)

³⁷² Hume continues 'Nothing can be more real, or concern us more, than our own sentiments' and that 'no more [than this] can be requisite to the regulation of our conduct and behaviour'.

³⁷³ Rabinowicz & Österberg (1996) (p. 12)

This might seem to be an ascription of something like Mackie's error theory to the subjectivist, but what is really ascribed to the subjectivist position here is, I believe, merely the second order, metaphysical error. The 'systematic misapprehension' referred to is thus not the one that gives the error theory its name, but rather the deeper misapprehension of the objective nature of value that non-error-theoretical subjectivists (such as Blackburn) also diagnoses. To sum up this line of reasoning: all subjectivists hold that insofar as people 'imagine a kind of objectivity for values' (to use Blackburn's words), they are guilty of a mistake. But this mistake is located at the second order level and concerns the nature of value. Whether this mistake leads to error also on the first order level is a further issue that is a matter of disagreement among subjectivists. This disagreement hangs on the nature of evaluative language, on which semantic view is held in conjunction with a subjectivist view about the nature of value.

Let me recapitulate the description of projectivism from the previous chapter, quoting Joyce again with my modifications:

Minimal projectivism:

- (1) 'We experience [value] as an objective feature of the world.'
- (2) 'This experience has its origin in some non-perceptual faculty; in particular, upon observing certain actions and characters (etc.) we have an affective attitude (e.g., the emotion of disapproval) that brings about the experience described in 1.'

As Joyce points out, the conjunction of 1 and 2 is both ontologically and semantically neutral. Nor does the conjunction of 1 and 2 entail 3 or my amended version 3'.

Metaphysical projectivism (in my amended form):

Minimal projectivism +

- (3') In fact, value is not an objective feature of the world.

As I argued in the previous chapter, even if they are not necessarily committed to any of its three component claims, it is natural for value subjectivists to accept metaphysical projectivism. Joyce then adds a fourth claim to metaphysical projectivism to yield error-theoretic projectivism.

Error-theoretic projectivism:

Metaphysical projectivism +

- (4) ‘When we utter sentences of the form “X is [valuable]” we are misdescribing the world; we are in error.’³⁷⁴

The conjunction of 1, 2 and 3 (or my 3’) does not, as Joyce points out, entail 4. If it did, non-cognitivist projectivism would be incoherent.

One way for a subjectivist to avoid the conclusions of the error theory whilst holding on to the diagnosis of metaphysical error—i.e. the error of thinking that value is an objective feature of the world—is to diagnose a *further* error. This further error lies in thinking that our evaluative language is descriptive. It surely seems to be the case that ordinary evaluative judgements are descriptive, but if that too is an illusion the first order error (captured in Joyce’s fourth claim) can be avoided. After all, if we are not describing anything in the first place we cannot be *mis*describing anything. This is the strategy of the non-cognitivist.

The non-cognitivist projectivist rejects the fourth claim of error-theoretic projectivism and replaces it with an alternative claim.

Non-cognitivist projectivism:

Metaphysical projectivism +

³⁷⁴ Joyce (2009) (p. 56). I continue to replace ‘morally wrong’ with ‘valuable’ to suit my purposes.

(4_{NC}) When we utter sentences of the form ‘X is good’ we are not (primarily) describing the world at all. Instead we are expressing some attitude towards X.

Joyce argues that whilst non-cognitivist projectivism is a coherent position, there might be an inherent tension in this view. As I have previously noted, Blackburn holds that it is ‘when we speak and think as though there were a property of things which our sayings describe, which we can reason about, know about, be wrong about, and so on’ that we project an attitude onto the world. But what is it to speak and think *as though* there were a property of things which our sayings describe? The natural answer, Joyce insists, is that

speaking as if the world were in a certain way amounts to *asserting* that it is that way, and thinking as if the world were in a certain way amounts to *believing* that it is that way. But asserting and believing are the hallmarks of a *cognitivist* attitude.³⁷⁵

One might perhaps suspect that Joyce’s formulation of metaphysical projectivism is somehow biased in favour of cognitivism and that a non-cognitivist might want to reformulate it to diminish the tension that Joyce points out. Let us first notice that saying that we experience value as an objective feature of the world is perhaps an overly theory-laden way to describe our phenomenology. Would a non-cognitivist like Blackburn really accept the first claim of Joyce’s projectivism? In a footnote, Blackburn says something that could make us doubt this: ‘I would deny that there is an aspect of moral phenomenology which gives morality an objective appearance which quasi-realism must regard as illusory’.³⁷⁶

Blackburn’s and many other non-cognitivist views deserve much closer consideration than I have been giving them here, but this goes beyond my present project. I conclude that, despite the tension identified by Joyce, non-cognitivism remains a viable option for value subjectivists.

³⁷⁵ Ibid (p. 58)

³⁷⁶ Blackburn (1985) (p. 21)

IV. Subjectivism and Error Theory

Mackie holds that a linguistic and conceptual analysis of our moral language would reveal a claim to objectivity, and if second order ethics was confined to such analyses then ‘it ought to conclude that moral values are at least objective: that they are so in part of what our ordinary moral statements mean’.³⁷⁷ Following Mackie, David Gauthier similarly says that ‘[i]f we were to suppose that the correct conception of value could be discovered by an analysis of ordinary language, we should no doubt be led to an objective conception (or perhaps to a conception with objective and subjective elements intertwined in hopeless confusion)’. Gauthier’s conclusion is that a plausible account of value must ‘include what Mackie calls an “error theory”’.³⁷⁸

‘If we reject the view that it is the function of [ethical] terms to introduce objective values into discourse about conduct and choices of action, there seem to be two main alternative types of account’³⁷⁹, Mackie writes. The two alternative accounts he has in mind are naturalism and non-cognitivism. Mackie believed that views of both these types capture part of the truth, but he insists that ‘both naturalist and non-cognitive analyses leave out the apparent authority of ethics’. Non-cognitive analyses do so by excluding ‘the claim to objective validity or truth’ whilst naturalist analyses do so by excluding ‘the categorically imperative aspect’ of evaluative terms and expressions.³⁸⁰ Mackie concludes: ‘Any analysis of the meanings of moral terms which omits [a] claim to objective, intrinsic, prescriptivity is to that extent incomplete; and this is true of any non-cognitive analysis, any naturalist one, and any combination of the two’.³⁸¹

I will not here take a stand on precisely how the error is to be characterized, or indeed whether there really is one. Like with semantically

³⁷⁷ Ibid

³⁷⁸ Gauthier (1986) (p. 58)

³⁷⁹ Mackie (1977) (p. 32)

³⁸⁰ Ibid (p. 33)

³⁸¹ Ibid (p. 35)

naturalistic views and non-cognitivist views, there is variation also among defenders of error theories. There are somewhat different analyses of the error that is thought to be inherent in evaluative language. Mackie sometimes characterized in terms of ‘objective prescriptivity’ and sometimes in terms of ‘objectively categorical imperatives’. Simon Kirchin has identified what he thinks of as a tension in error theories:

In order to make their position convincing moral error theorists must alight on a particular idea and argue that it is a crucial commitment of everyday moral thought and language. Further, in order to be sure that they can convict a commitment of error, error theorists might need to specify that commitment so as to rule out alternative defensible understandings. The danger is that the more one specifies the formulation of a commitment, the more it is likely that the commitment is less than crucial to many people’s everyday moral thought and language. . . . Hence, it is questionable whether error theorists can plausibly convict everyday moral thought and language overall of error.³⁸²

I will not try to resolve this alleged tension by offering a detailed analysis of evaluative language or the claims inherent in such language. Instead, I will merely point out that insofar as it is a crucial commitment of everyday evaluative thought and language that values are not grounded on attitudes, then value subjectivism as I have described it leads to an error theory. If, when it is thought and asserted that ‘X is good’, X is thought and claimed not merely to be good-relative-to some first-person perspective, but good from the third-person perspective, then the thought and assertion do seem to embody or include a mistake.

As Mackie notes, since his error theory ‘goes against assumptions ingrained in our thought and built into some of the ways in which language is used, since it conflicts with what is sometimes called common sense, it needs very solid support. . . . If we are to adopt this view, we must argue explicitly for it’.³⁸³ Mackie then goes on to provide arguments.³⁸⁴ But

³⁸² Kirchin (2010) (p. 167)

³⁸³ Mackie (1977) (p. 35)

³⁸⁴ *Ibid* (pp. 36–49). I think that the argument from relativity and the various strands of the argument from queerness are best seen as sub-parts of an overall argument that also includes

whether his much discussed arguments (or the arguments of other error theorists) deliver the solid support needed is an issue that lies beyond the present inquiry.³⁸⁵

Note on Truth and Knowledge

Subjectivism, I have argued, implies that from the third-person perspective things are neither good nor bad, neither better nor worse—nothing matters from the point of view of the universe. It is only from the first-person perspectives of individual subjects that things can be good or bad. If we hold that truth is an inherently objective notion, that truth is essentially objective truth, then if something is true it is true from the third-person point of view and not merely ‘true’ from this or that particular first-person perspective. Value subjectivism, at least when combined with an error theory of evaluative language, therefore implies that there are no evaluative truths.³⁸⁶ This, in turn—given the classic, three-part definition of knowledge as justified, true belief—entails that there is no evaluative knowledge. It cannot be known that something is good or bad. I have argued that there are objective truths about what is good or bad relative to this or that person, but

Mackie’s projectivism or objectification theory as an integral part. Richard Garner (forthcoming) stresses the inference-to-the-best-explanation nature of Mackie’s overall argument.

³⁸⁵ An interesting recent proposal by Toby Svoboda (2011) is a hybrid theory that combines elements of an error-theory with elements of an expressivist meta-ethical theory. Svoboda argues that such a hybrid is theoretically preferable over pure error-theory as well over pure expressivism. One potential benefit with adopting such a hybrid theory (that comes in virtue of its expressivist element) is that it would allow the subjectivist to subscribe to judgement internalism. The same would be the case if the subjectivist adopted a pure expressivist view, but then other benefits of the hybrid view (those brought by the cognitivist elements) would be lost. To the latter, Svoboda counts the ability to handle the so-called Frege-Geach problem (pp. 41–43).

³⁸⁶ However, a subjectivist who buys into the quasi-realist programme could treat sentences like ‘It is true that X is good’, ‘The proposition “X is good” is true’, and ‘I know that X is good’, etc. as saying no more than ‘X is good’, which, in turn, might perhaps be given a non-cognitivist analysis as doing no more than expressing a positive attitude towards X. I will not here evaluate the prospects of this strategy.

such truths are not *evaluative* truths. Hence, even though such good-relative-to truths can be known, that would not constitute evaluative knowledge. Knowing that X is good-relative-to-S does not entail that we know that X is good or that X is good *for* anyone. A lot more would need to be said about this, but it does seem that value subjectivism naturally leads to scepticism about evaluative truth and knowledge.

There may be a class of statements involving typical value words such as ‘good’ and ‘bad’ that everyone can agree to be objectively true independently of the subjectivism versus objectivism debate. Mackie acknowledges this: ‘[T]here are certain kinds of value statements which undoubtedly can be true or false, even if, in the sense I intend, there are no objective values.’³⁸⁷ His examples all include evaluations made in relation to assumed and agreed upon standards, such as the classing of wool, the grading of apples and the awarding of prizes at sheepdog trials. Given sufficiently determinate standards, which may be written down or just tacitly understood, ‘it will be an objective issue, a matter of truth and falsehood, how well any particular specimen measures up to those standards. Comparative judgements [relative to standards] in particular will be capable of truth and falsehood: it will be a factual question whether this sheepdog has performed better than that one.’³⁸⁸ The subjectivist about values³⁸⁹, Mackie argues, does not deny that there can be objective evaluations according to standards and that these are as possible in the aesthetic as in the moral fields: ‘justice or injustice of decisions relative to standards can be a thoroughly objective matter’ (‘though there may still be a subjective element in the interpretation or application of standards’). But the objectivity involved here is only the objectivity of the relation of satisfying a certain standard, not the objectivity of value.³⁹⁰ Evaluation of this kind doesn’t seem

³⁸⁷ Mackie (1977) (p. 25)

³⁸⁸ Ibid (p. 26)

³⁸⁹ Mackie (1977) introduces the name *subjectivism* as one out of two alternative labels for the second order thesis that there are no objective values (p. 15). Beyond the general remark that ‘if values are not objective they are in some very broad sense subjective’ (p. 18), Mackie’s subjectivism is a negative doctrine in the sense that ‘it says what there isn’t, not what there is’ (p. 17). I have attempted a more positive characterization of value subjectivism by focusing on the thesis that values are grounded on attitudes.

to be what subjectivists and objectivists disagree on or seek to analyse in different ways.

Some of Mackie's examples of value statements that can be objectively true even if there are no objective values concern virtues. He says that it is undeniable that there is a difference between a brave man and a coward, and between a kind action and a cruel one, and that this is an objective matter. But, he insists, this is so in virtue of the natural, descriptive differences between these men and actions, not their differences in value.

It is a hard fact that cruel actions differ from kind ones, and hence that we can learn, as in fact we all do, to distinguish them fairly well in practice, and to use the words 'cruel' and 'kind' with fairly clear descriptive meanings; but is it an equally hard fact that actions which are cruel in such a descriptive sense are to be condemned?³⁹¹

The debate between subjectivists and objectivists concerns the objectivity specifically of value, not the objectivity of the natural, factual, differences on the basis of which differing values are ascribed—the features upon which values supervene. Subjectivists and objectivists can agree that there are descriptive constraints on how words like 'brave' and 'kind' can be applied, and insofar as these descriptive elements are concerned nobody need deny the possibility of plain truth. It would be a mistake, an objective falsehood, to call someone who systematically runs away from danger 'brave', or someone who consistently faces danger 'a coward'. But this would be so in virtue of the descriptive components of such thick value concepts as *courage* and *cowardice*.³⁹² It would be a factual mistake—a mistake about ordinary,

³⁹⁰ Mackie (*ibid*) also points out that the subjectivist is not committed to the implausible claim that the choice of basic standards is therefore an arbitrary matter. On the contrary, Mackie insists, 'the standards used in sheepdog trials clearly bear some relation to the work that sheepdogs are kept to do, the standards for grading apples bear some relation to what people generally want in or like about apples, and so on' (p. 27).

³⁹¹ *Ibid* (p. 17)

³⁹² There may be disagreements over the precise descriptive content of thick value concepts, and over the extent to which the evaluative/normative and descriptive elements of thick value concepts can be separated. Insofar as this is a problem it is a problem for subjectivists and objectivists alike. Michael Smith (2013) has questioned whether there really are any purely

non-evaluative facts and properties—not a mistake about some peculiarly evaluative or normative facts and properties.

thin value concepts and argued that ‘the very feature of thick ethical concepts that explains the non-ethical constraint on their application is a feature of thin ethical concepts as well. The upshot is that all ethical concepts, even the maximally thin ones, turn out to be a little bit thick’ (p. 98).

Afterword

Two separate but interrelated projects have been running in tandem throughout this work. One project has been to map out different possibilities: different routes for a subjectivist to take. On numerous occasions I have pointed out that beyond subjectivism's core claim that value is grounded on attitudes, different subjectivists may have different views, and depending on which particular views are adopted, we can speak of different forms or variants of value subjectivism. In this regard I cannot lay a claim to have been exhaustive; there are probably possibilities I have overlooked. The other project has been to identify which of the several different possible variants of subjectivism is the most plausible one. Even though I have left some internal debates unresolved, I hope that at least in outline a plausible candidate theory has emerged. In this brief afterword I will recount some of the central features of the subjectivist analysis of value and also raise some issues that remain for future investigation.

I have argued that the central claim of the value subjectivist is that value is grounded on attitudes. This claim does not, contrary to common assumption, imply that the objects of value need themselves involve attitudes. It is generally not the case that what makes something good is that it is desired or preferred by someone. In particular, subjectivism does not imply that only preference satisfaction is valuable for its own sake. Instead, value is by the subjectivist taken to accrue to the objects of our attitudes in virtue of the features of these objects for which we favour or disfavour them. Given the empirical fact of the heterogeneity of our attitudes this naturally leads to a pluralistic axiology on which many different things carry final value. Such a pluralistic axiology is, I believe, more in line with common sense than monistic axiologies according to which one and only one thing is valuable for its own sake.

Some philosophers have held that only objectivists can speak about intrinsic value or that in order for subjectivists to do so they would need to rethink the notion. I have argued that, like objectivists, subjectivists too can

talk about intrinsic final values, meaning by *intrinsic* the same thing as objectivists: namely, a value that supervenes solely on the internal features of its bearer. On the subjectivist analysis, an object is finally valuable when it is favoured for its own sake and (finally) intrinsically valuable when the properties for which it is favoured (for its own sake) are all internal to the object. The attitude on which the value is grounded is not itself counted among the features upon which the intrinsic value supervenes.

This analysis distinguishes between the supervenience base of value, the features of the object upon which its value supervenes, and the constitutive ground of the value, the attitude that the subject is directing towards the object. This distinction has been of central importance in my project of trying to understand the subjectivist position. Not allowing the subjectivist to draw the distinction between supervenience base and constitutive ground would, I argued, be to beg the question against the subjectivist. Insisting that the attitude (that the subjectivist thinks of as the constitutive ground) should really be treated as itself being part of the supervenience base would not only force upon us an overly restrictive view of value bearers—that the only bearers of value are states of affairs involving some subject having some attitude towards something—it would also force us to say that such states have their values regardless of whether *they* are objects of attitudes. This would be asking subjectivists to accept objective values.

The availability of this analysis of intrinsic value shows that there is no need for subjectivists to concede a monopoly of analysing intrinsic value to objectivists. Nor is there any need for subjectivists to rethink the traditional notion of intrinsic value in order to accommodate it into a subjectivist framework. That subjectivists and objectivists should be seen as offering rival analyses of the same thing has been treated as a guiding assumption throughout this work. Holding on to the traditional notion of intrinsic value whilst recognizing that subjectivists can analyse intrinsic value in the same traditional sense straightforwardly allows us to maintain the assumption that subjectivism and objectivism are rival positions.

Subjectivism has been accused of implying that if we had different attitudes than those that we actually have then different things would have been valuable. If people had desired a paradigmatically worthless thing such as a saucer of mud for its own sake, and lacked any pro attitude towards paradigmatically good things such as knowledge, peace and friendship, for example, would not the saucer of mud be finally valuable and knowledge, peace and friendship would lack value? I have argued that this does not follow since it is our actual attitudes that ground values even in hypothetical

situations. The attitudes that we could have had but we do not actually have enter the picture only insofar as our actual attitudes are sensitive to our hypothetical attitudes.

Does subjectivism imply an objectionable form of relativism? I have argued that even though subjectivism is inherently relativistic in one sense, it does not imply that there need be anything relativistic in the contents of our evaluations. Furthermore, though things can only be good or bad relative to someone it does not follow that things can be good or bad only by being good or bad *for* someone. Subjectivists can recognise non-relational values such as goodness *simpliciter* as well as relational values such as goodness-for. I sketched the outlines of a subjectivist analysis of goodness-for that involves attitudes of favouring something for someone's sake. Roughly, something is good for some person relative to a subject S insofar as it is favoured by S for that person's sake. Whether this is, as it stands, the most plausible subjectivist analysis of goodness-for remains an open question, but the availability of this analysis is enough to show that it is not impossible for subjectivists to analyse relational values.

On the individualist version of subjectivism, the relevant value-grounding attitudes are those of the individual subject. Other people's attitudes enter the picture only insofar as the relevant subject's attitudes are sensitive to their attitudes. Does this imply that this form of value subjectivism is objectionably self-regarding or egoistic? I have argued that subjectivism does not have this implication because the individual subject's other-regarding or altruistic attitudes ground values just as well as his or her self-regarding or egoistic attitudes do. There is no reason to believe that all our attitudes are self-regarding, though undoubtedly a few of them are.

The basic components of the subjectivist analysis of the nature of value include subjects having attitudes towards objects in virtue of non-evaluative features of these objects. Concerning each of these elements of the basic analysis one can ask further questions. What or who qualifies as a subject? A provisional answer may be that whatever has attitudes count as a subject in the relevant sense. But which, then, are the relevant attitudes? Some subjectivists may say that it is preferences, others that it is desires and yet others would want to include a much broader set of pro and con attitudes. I have argued that the relevant attitudes are specifically conative attitudes, attitudes that involve a motivational element. That attitudes possess a positive/negative polarity or valence is something that I have taken for granted. But questions may remain concerning the ultimate nature of this polarity. What is it that makes a positive attitude positive and a negative

attitude negative? This is one of the questions that remain for future investigation.

Yet another question concerns what kinds of objects can be bearers of value. Here the answer I have given is that anything that may be an object of a pro or con attitude may be a bearer of value. If the relevant attitudes are thought to be preferences, then it is perhaps natural to think of the objects of value as being states of affairs, but if it is held that a broader set of attitudes constitute values then it seems more natural to be a value-bearer pluralist and say that objects belonging to several different ontological categories may be bearers of value.

There are also questions concerning idealization. I have argued that it is the attitudes that we actually have here and now that ground values and not the attitudes that we would have in some idealized circumstance. Yet, if an object of an actual, present attitude lacks some or all of the features for which it is favoured by the subject, the attitude does not confer value on that object. Instead, this attitude may confer value on other objects that in fact do possess the relevant features even though these objects have never been the objects of any attitude. Our actual, present attitudes may be subjected to further restrictions consistent with subjectivism. Plausible candidates include formal constraints of instrumental rationality arising from the practice of pursuing what is valuable.

In all of these areas there is room for variation among value subjectivists. Add to this that subjectivists may have different views about the ontological status of value—realism, anti-realism, or constructivism, though I have argued that anti-realism is the most natural view for a subjectivist—and the semantics of evaluative language—cognitivist and non-cognitivist—it becomes clear that we arrive at a wide variety of views that all agree on the central claim that value is grounded on attitudes, but that beyond this core they may differ in numerous interesting ways. The subjectivist family may be as diverse as the objectivist family.

Even though I do think that we should accept a subjectivist theory of value, I have not argued for that conclusion in this work. Doing that would require undermining objectivist positions, either by some general anti-objectivist argument or by attacking individual objectivist positions on a case-by-case basis. My own subjectivist stand is based on the belief that subjectivism provides a better explanation of observed phenomena than does objectivism, but making good on this claim is not the project that I have undertaken here. I have focused on what value subjectivism is and its implications. However, I do hope that the considerations that I have raised

have brought value subjectivism into a somewhat more favourable light than that in which it is sometimes held.

Bibliography

- Anscombe, Elisabeth (2000/1957). *Intention* (2nd edn.). Harvard University Press
- Beardsley, Monroe C. (1965). 'Intrinsic Value'. *Philosophy and Phenomenological Research*, 26: 1-17
- Blackburn, Simon (1971). 'Moral Realism' in Casey, John (ed.), *Morality and Moral Reasoning*. Methuen & Co. Ltd.
- (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Clarendon Press
- (1985). 'Errors and the Phenomenology of Value' in Honderich, Ted (ed.) *Morality and Objectivity: A Tribute to J. L. Mackie*. Routledge & Kegan Paul
- (1988). 'How to be an Ethical Antirealist'. *Midwest Studies in Philosophy*, 12: 361-375
- Brown, Donald E. (1991). *Human Universals*. McGrath-Hill. Updated version of the list published as an appendix in Pinker, Steven (2002). *The Blank Slate: The Modern Denial of Human Nature*. Penguin Books Ltd.
- Brännmark, Johan (2002). *Morality and the Pursuit of Happiness: A Study in Kantian Ethics*. Media-Tryck, Lund University
- Butler, Joseph (1900/1729). 'Fifteen Sermons Preached at the Rolls Chapel' in *The Works of Bishop Butler*, (vol. 1). Macmillan
- Bykvist, Krister (1996). 'Utilitarian Deontologies?: On Preference Utilitarianism and Agent-Relative Value' in Rabinowicz, Wlodek (ed.). *Preference and Value: Preferentialism in Ethics*. Studies in Philosophy, 1996:1, Department of Philosophy, Lund University. (Reprinted from *Theoria*, 62: 124-143)
- Copp, David & Sobel, David (2001). 'Against Direction of Fit Accounts of Belief and Desire'. *Analysis*, 61(1): 44-53
- Crisp, Roger (2000). Review of *Value ... and What Follows* by Joel Kupperman, *Philosophy*, 75: 458-462
- Dancy, Jonathan (2000). *Practical Reality*. Oxford University Press
- (2004). *Ethics without Principles*. Oxford University Press

- Darwall, Stephen (1983). *Impartial Reason*. Cornell University Press
- (1992). 'Internalism and Agency'. *Philosophical Perspectives*, 6: 155-174
- (2000). 'Normativity and Projection in Hobbes's Leviathan'. *The Philosophical Review*, 109(3): 313-347
- Double, Richard (2006). *Metaethical Subjectivism*. Ashgate
- Enoch, David (2005). 'Why Idealize?' *Ethics*, 115: 759-787
- Garner, Richard (work in progress). 'A Unified Argument for the Moral Error Theory'. (Draft published by the author online at <http://beyondmorality.com> 2013-10-16)
- Gauthier, David (1986). *Morals by Agreement*. Oxford University Press
- Geach, Peter (1956). 'Good and Evil'. *Analysis*, 17: 33-42
- Hare, Richard M. (1981). *Moral Thinking: Its Level, Method and Point*. Clarendon Press
- (1985). 'Ontology in Ethics' in Honderich, Ted (ed.) *Morality and Objectivity: A Tribute to J.L. Mackie*. Routledge & Kegan Paul
- Harman, Gilbert (1977). *The Nature of Morality: An Introduction to Ethics*. Oxford University Press
- Harvey, Martin (2004). 'Hobbes and the Value of Justice'. *The Southern Journal of Philosophy*, 42: 439-452
- Hobbes, Thomas (1991/1658). *De Homine*. (Published as one volume together with *De Cive* as Gert, Bernard (ed.) *Man and Citizen*) Hackett
- (1996/1651). *Leviathan*. Oxford University Press
- Huemer, Michael (1996). 'The Subjectivist's Dilemma'. *Objectivity*, 2(4): 77-92. (Available on the author's website: <http://home.sprynet.com/~owl1/subj.htm>)
- Hume, David (1978/1739-40). *A Treatise of Human Nature*. Selby-Bigge, L. A. (ed.). 2nd edn. revised by Nidditch, P. H. Oxford University Press
- (1983/1751). *An Enquiry Concerning the Principles of Morals*. J. B. Schneewind (ed.). Hackett
- Johnston, Mark (1989). 'Dispositional Theories of Value'. *Proceedings of the Aristotelian Society (supplementary volumes)*, 63: 139-174
- Joyce, Richard (2006). *The Evolution of Morality*. MIT Press
- (2009). 'Is Moral Projectivism Empirically Tractable?' *Ethical Theory and Moral Practice*, 12: 53-75

- (2010). ‘Patterns of Objectification’ in Joyce, Richard & Kirchin, Simon (eds.) *A World Without Values: Essays on John Mackie’s Moral Error Theory*. Springer Press
- (2012). ‘Metaethical Pluralism: How Both Moral Naturalism and Moral Scepticism May be Permissible Positions’ in Nuccetelli, Susana & Seay, Gary (eds.) *Ethical Naturalism: Current Debates*. Cambridge University Press. (Page references in notes to the penultimate draft published online by the author at http://personal.victoria.ac.nz/richard_joyce/onlinepapers.html)
- Kagan, Shelley (1992). ‘The Limits of Well-Being’ in Frankel Paul, Ellen, Miller, Fred D. & Paul, Jeffrey (eds.) *The Good Life and the Human Good*. Cambridge University Press
- (1998). ‘Rethinking Intrinsic Value’. *The Journal of Ethics*, 2: 277-297
- Kant, Immanuel. (1774/1790). *Kritik der Urteilkraft*. Suhrkamp Verlag
- Keller, Simon. (2009). ‘Welfarism’. *Philosophy Compass*, 4(1): 82-95
- Korsgaard, Christine M. (1983). ‘Two Distinctions in Goodness’. *The Philosophical Review*. 92: 169-195. Reprinted in Korsgaard, Christine M. (1996). *Creating the Kingdom of Ends*. Cambridge University Press
- (2013) ‘The Relational Nature of the Good’. *Oxford Studies in Metaethics*, (vol. 8). Oxford University Press
- (2014) “On Having a Good”. *Philosophy*, 89(3): 405-429
- Kraut, Richard. (2011). *Against Absolute Goodness*. Oxford University Press
- Kripke, Saul (1980). *Naming and Necessity*. Harvard University Press
- Lewis, David (1989). ‘Dispositional Theories of Value’. *Proceedings of the Aristotelian Society, (supplementary volumes)*, 63: 113-137
- Loeb, Don (2010). ‘The Argument from Moral Experience’ in Joyce, Richard & Kirchin, Simon (eds.) *A World Without Values: Essays on John Mackie’s Moral Error Theory*. Springer Press
- Locke, John (1775/1689). *An Essay Concerning Human Understanding*. Oxford University Press
- Mack, Eric (1993). ‘Agent-Relativity of Value, Deontic Restraints, and Self-Ownership’ in Frey, R. G. & Morris, Christopher W. (eds.) *Value, Welfare, and Morality*. Cambridge University Press
- Mackie, John L. (1946). ‘A Refutation of Morals’. *Australasian Journal of Psychology and Philosophy*, 24: 77-90
- (1977). *Ethics: Inventing Right and Wrong*. Penguin Books

- (1980). *Hume's Moral Theory*. Routledge
- Mill, John Stuart (2001/1864). *Utilitarianism*. (2nd edn.) Hackett
- Moore, Adam D. (2004) 'Values, Objectivity, and Relationalism'. *The Journal of Value Inquiry*, 38: 75-90
- Moore, G. E. (1993/1903). *Principia Ethica*. (Revised edn.) Cambridge University Press
- (1922). 'The Conception of Intrinsic Value' in *Philosophical Studies*. Routledge & Kegan Paul
- Nagel, Thomas (1986). *The View from Nowhere*. Oxford University Press
- O'Neill, John (1992). 'The Varieties of Intrinsic Value'. *Monist*, 75: 119-137
- Parfit, Derek (1984). *Reasons and Persons*. Oxford University Press
- (2011). *On What Matters (vol. 1)*. Oxford University Press
- Perry, Ralph Barton (1950/1926). *The General Theory of Value*. Harvard University Press
- Persson, Ingmar (2005). *The Retreat of Reason*. Oxford University Press
- Pettit, Phillip & Smith, Michael (1990). 'Backgrounding Desire'. *The Philosophical Review*, 99: 565-592
- Pufendorf, Samuel (2011/1672). *The Law of Nature and of Nations*. Nabu Press
- Rabinowicz, Wlodek (1996). 'Värde Grundat på Preferens: Om Två Tolkningar av Preferensutilitarism', *Filosofisk Tidskrift*, 4: 21-36
- Rabinowicz, Wlodek & Österberg, Jan (1996). 'Value Based on Preferences: On Two Interpretations of Preference Utilitarianism'. *Economics and Philosophy*, 12: 1-27. (Reprinted in Rabinowicz, Wlodek (ed.). *Preference and Value: Preferentialism in Ethics*. Studies in Philosophy 1996:1, Department of Philosophy, Lund University)
- Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (2000). 'A Distinction in Value: Intrinsic and for its Own Sake'. *Proceedings of the Aristotelian Society*, 100(1): 33-51
- (2004). 'The Strike of the Demon: On Fitting Pro-Attitudes and Value'. *Ethics*, 114: 391-423
- Rachels, James (2007). *The Elements of Moral Philosophy* (5th edition). McGraw-Hill
- Ripstein, Arthur (1993). 'Preference' in Frey, R. G. & Morris, Christopher W. (eds.) *Value, Welfare, and Morality*. Cambridge University Press. Reprinted (with a new conclusion) in Morris, Christopher W. & Ripstein, Arthur (eds.)

- (2001). *Practical Rationality and Preference: Essays for David Gauthier*. Cambridge University Press
- Rosati, Connie (2008). 'Objectivism and Relational Good'. *Social Philosophy & Policy Foundation*, 25(1): 314-349
- Ross, W. D. (1930). *The Right and the Good*. Clarendon Press
- Russell, Bertrand. (1984). 'Is There an Absolute Good?' *The Journal of Bertrand Russell Studies*: 6(2):144-149. (Believed to have been written in 1922. First published in 1984 with an introductory note by Alan Ryan.)
- Rønnow-Rasmussen, Toni (2002). 'Hedonism, Preferentialism, and Value Bearers'. *Journal of Value Inquiry*, 36: 463-472
- (2003). 'Subjectivism and Objectivism' in Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (eds.) *Patterns of Value: Essays on Formal Axiology and Value Analysis (vol. 1)*. Lund Philosophy Reports 2003:1, Lund University
- (2007). 'Analysing Personal Value'. *The Journal of Ethics*, 11: 405-435
- (2009). 'On For Someone's Sake Attitudes'. *Ethical Theory and Moral Practice*, 12: 397-411
- (2011). *Personal Value*. Oxford University Press
- (2013). 'Good and Good For'. *The International Encyclopedia of Ethics*
- Sainsbury, R. M. (1998). 'Projections and Relations'. *The Monist*, 81(1): 133-160
- Scanlon, Thomas M. (1998). *What We Owe to Each Other*. Harvard University Press
- Scheffler, Samuel (1987). *The Rejection of Consequentialism: A Philosophical Investigation of Considerations Underlying Rival Moral Conceptions*. Clarendon Press
- Schroeder, Mark (2007a). 'Teleology, Agent-Relative Value, and "Good"'. *Ethics*, 117: 265-295
- (2007b). *Slaves of the Passions*. Oxford University Press
- Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford University Press
- Sidgwick, Henry (1981/1907). *The Methods of Ethics* (7th edn.). Hackett
- Smith, James Ward (1948). 'Senses of Subjectivism in Value Theory'. *The Journal of Philosophy*, 45(15): 393-405
- Smith, Michael (2003). 'Neutral and Relative Value after Moore'. *Ethics*, 113: 576-598

- (2013). ‘On the Nature and Significance of the Distinction between Thick and Thin Ethical Concepts’ in Kirchin, Simon (ed.) *Thick Concepts*. Oxford University Press
- Sobel, David (2009). ‘Subjectivism and Idealization’. *Ethics*, 119: 336-352
- Svoboda, Toby. (2011). ‘Hybridizing Moral Expressivism and Moral Error Theory’. *The Journal of Value Inquiry*, 45: 37-48
- Thomson, Judith Jarvis (1997). ‘The Right and the Good’. *The Journal of Philosophy*, 94(6): 273-298
- Williams, Bernard (1981). ‘Internal and External Reasons’ in Williams, Bernard (1981). *Moral Luck*. Cambridge University Press. (Reprinted from Harrison, Ross (ed.) *Rational Action*. Cambridge University Press)
- Zimmerman, Michael J. (2001). *The Nature of Intrinsic Value*. Rowman & Littlefield