



DEREK PARFIT

On What Matters

VOLUME THREE

On What Matters

On What Matters

VOLUME THREE

DEREK PARFIT

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trademark of
Oxford University Press in the UK and in certain other countries

© Derek Parfit 2017

The moral rights of the author have been asserted

First Edition published in 2017

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence, or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2009029662

ISBN 978-0-19-877860-8

Printed in Great Britain by
Clays Ltd, St Ives plc

Links to third-party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third-party website referenced in this work.

To
Larry Temkin, Jeff McMahan, and Ruth Chang

On What Matters

VOLUME ONE

List of Contents

Introduction

Preface

Summary

PART ONE Reasons

PART TWO Principles

PART THREE Theories

APPENDICES

Notes to Volume One

References

Bibliography

Index

VOLUME TWO

List of Contents

Preface

Summary

PART FOUR Commentaries

PART FIVE Responses

PART SIX Normativity

APPENDICES

Notes to Volume Two

References

Bibliography

VOLUME THREE

List of Contents

Preface

Summary

PART SEVEN Irreducibly Normative Truths

PART EIGHT Expressivist Truths

PART NINE Normative and Psychological Reasons

PART TEN Ethics

Notes to Volume Three

References

Bibliography

Index

Contents

VOLUME THREE

PREFACE	xiii
SUMMARY	1

PART SEVEN

IRREDUCIBLY NORMATIVE TRUTHS

37	HOW THINGS MIGHT MATTER	41
128	Caring and Having Reasons to Care	41
129	Philosophical Disagreements	50
38	NON-REALIST COGNITIVISM	55
130	Meta-Ethics	55
131	Ontology	58
39	NORMATIVE AND NATURAL TRUTHS	65
132	Concepts and Properties	65
133	The Co-Extensiveness Argument	71
134	The Normativity Objection	72
135	Scientific Analogies	75
136	The Triviality Objection	84
40	GIBBARD'S OFFER TO NON-NATURALISTS	87
137	The Single Property Illusion	87
138	Naturalistic States of Affairs and Normative Truths	88

41	RAILTON'S DEFENCE OF SOFT NATURALISM	91
139	The Identity of Properties	91
140	Railton's First Response to the Triviality Objection	95
42	RAILTON'S RESOLUTION OF OUR DISAGREEMENTS	99
141	Railton's Wider View	99
142	What is Achieved by Railton's Wider View	102
143	Railton's Commentary	113
43	JACKSON'S NON-EMPIRICAL NORMATIVE TRUTHS	128
144	Jackson's Co-Extensiveness Argument	128
145	Jackson's Metaphysical Assumptions	132
44	SCHROEDER'S CONSERVATIVE REDUCTIVE THESIS	137
146	Schroeder's Criticisms of the Triviality Objection	137
147	How Schroeder and I Could Resolve our Disagreements	143
148	How I Misled Russell	159

PART EIGHT

EXPRESSIVIST TRUTHS

45	QUASI-REALIST EXPRESSIVISM	165
149	Desires, Attitudes, and Beliefs	165
150	Earning the Right to Talk of Moral Truth	176
46	GIBBARD'S RESOLUTION OF OUR DISAGREEMENTS	182
151	Gibbard's Convergence Claim	182
152	Does it Matter Whether Things Matter?	184
153	Getting It Right	191
154	Gibbard's Metaphysical Doubts	197

47	ANOTHER TRIPLE THEORY	205
155	Gibbard's Commentary	205
156	A Happy Ending	224

PART NINE

NORMATIVE AND PSYCHOLOGICAL REASONS

48	EXPRESSIVIST REASONS	241
157	Blackburn's Bafflement	241
158	Blackburn's Beliefs about Reasons	244
49	SUBJECTIVIST REASONS	251
159	Smith's Defence of Subjectivism	251
160	Street's Defence of Subjectivism	255
50	STREET'S META-ETHICAL CONSTRUCTIVISM	264
161	Street's Debunking Arguments	264
162	Street's Relativism	268
163	The Normative Implications of Street's View	279
164	Chappell's Claims about Street	285
51	MORALITY, BLAME, AND INTERNAL REASONS	291
165	Darwall's Claims about Internal Reasons	291
166	Darwall's Defence of Moral Internalism	299
52	NIETZSCHE'S MOUNTAIN	309
167	Nietzsche and the Convergence Claim	309

PART TEN

ETHICS

53	WHAT MATTERS AND UNIVERSAL REASONS	315
168	The Any-All Thesis	315
169	Universalism about What Matters	330

54	CONFLICTING REASONS	335
170	Sidgwick's Problem	335
171	Moral and Self-interested Reasons	338
172	Other Problems	342
55	THE RIGHT AND THE GOOD	349
173	Moral Ambivalence	349
174	The Badness of Wrongdoing	354
175	Moral and Impartial Reasons	357
176	Wrongness and Reasons	366
56	DEONTOLOGICAL PRINCIPLES	369
177	The Means Principle	369
178	Harming and Saving from Harm	383
57	ACT CONSEQUENTIALISM AND COMMON SENSE MORALITY	395
179	Good and Bad Ways of Treating People	395
180	Deontic and Non-Deontic Badness	400
181	Personal and Shared Duties	406
58	TOWARDS A UNIFIED THEORY	413
182	Act Consequentialism	413
183	Rule and Motive Consequentialism	417
184	Optimific Motives and Rules	419
185	Small Effects and Great Harms	422
186	Inconclusive Conclusions	433
	<i>End Notes</i>	439
	<i>References</i>	451
	<i>Bibliography</i>	461
	<i>Index</i>	465

Preface

Though this book is called *Volume Three*, it can be read on its own. The start of the Index explains how you could read even less.

I am very grateful to Peter Singer, without whom I would have written none of *Volume Three*. Singer persuaded some very good philosophers to write the papers in *Does Anything Really Matter?*, the companion volume to this book. I apologize to the writers of these papers for taking so long to write my responses to them. These papers showed me that I had made some bad mistakes, and led me to have some new ideas. I was also thrilled to discover that two of the writers of these papers, Peter Railton and Allan Gibbard, had independently suggested how we might be able to resolve at least some of our main meta-ethical disagreements. I defend these suggestions in Chapters 38, 39, 40, 42, 46, and 47. In their commentaries, which are included in these chapters, Railton agrees that our disagreements have been wholly resolved, and Gibbard agrees that our disagreements have been partly resolved. I am deeply worried by disagreements with people who seem as likely as I am to be getting things right. That is why, like Railton, I find it ‘immensely heartening’ that Railton, Gibbard, and I now have similar beliefs.

Singer also made a remark which led me to write the rest of this book. He politely expressed his disappointment that, in my *Volumes One* and *Two*, I say little about the disagreements between Act Consequentialists and believers in what Sidgwick called Common Sense Morality. In Part Ten of this book one of my aims is to show that some of these disagreements can be resolved. As I wrote some years ago: ‘Non-Religious Ethics is at a very early stage. We cannot yet predict whether . . . we will all reach agreement. Since we cannot know how Ethics will develop, it is not irrational to have high hopes.’

I have been helped by many other people. I have been helped most by Selim Berker, Ruth Chang, Frances Kamm, Jeff McMahan, Ingmar Persson, Tim Scanlon, Sharon Street, and Larry Temkin. I have been greatly helped by Robert Audi, John Broome, Nicholas Bostrom, Roger Crisp, Garrett Cullity, Jonathan Dancy, David Enoch, William Fitzpatrick, Thomas Hurka, Thomas Nagel, Michael Otsuka, Samuel Scheffler, and Knut Skarsaune. Other people who have helped me are Marcello Antosh, Benjamin Butler, David Copp, Andrew Forcehimes, Daniel Forman, James Goodrich, Adil Ahmed Haque, Andrew Harris, Christopher Hauser, Hasan Dindjer, Frank Jackson, Aaron Jaslove, Guy Kahane, Justin Kalef, Joseph Kerstein, Douglas Kremm, Anton Markoc, Daniel Munoz, Jake Nebel, Martin O'Neill, Toby Ord, Jacob Ross, Richard Rowland, Bruce Russell, Bart Schultz, Kieran Setiya, John Skorupski, Saul Smilansky, Sigrun Svavarsdottir, Victor Tadros, Fiona Woollard, Alex Worsnip, Frank Wu, and, I am sure, several other people whose names I either failed to write down or cannot find.

I am also very grateful to Peter Momtchiloff for giving me, yet again, much wise advice.

SUMMARY

PART SEVEN IRREDUCIBLY NORMATIVE TRUTHS

CHAPTER 37 HOW THINGS MIGHT MATTER

128 *Caring and Having Reasons to Care*

When we claim that some things matter, we might mean only that these things matter to people. Suffering matters, for example, in the sense that people care about suffering. No one doubts that some things matter in this psychological sense. Some things also matter, I believe, in the different, normative sense that we have reasons to care about these things.

Gibbard believes that some things matter in a third, expressivist sense. When we say that suffering matters, Gibbard claims, we are telling people to care about suffering. It is unclear how things might matter in this expressivist sense. Gibbard would not claim that something matters when and because he tells us to care about this thing. But Gibbard also believes that, in saying that suffering matters, he is getting it right. Since most of us would have similar beliefs, Gibbard could include such beliefs in his account of what we mean. When we say that suffering matters, Gibbard could claim, we are both telling people to care about suffering, and claiming that, in telling people to care, we are getting it right. If we *are* getting it right, this claim would be true. Suffering would matter in this wider expressivist sense.

Temkin suggests that, to refute the Nihilistic view that nothing matters, it is enough to point out that some things matter to people. But Nihilists don't mean that no one cares about anything. Nihilists mean

that nothing matters in any significant normative sense. Temkin also suggests that, even if nothing mattered in the normative, reason-implying sense, there are some weaker normative senses in which some things would matter. That is not, I claim, true.

129 *Philosophical Disagreements*

Though Temkin and I both believe that some things matter in the reason-implying sense, and our views are in other ways similar, Temkin claims that I should not have wholly rejected views that conflict with ours, but should have looked for points of agreement with these other views. In writing this book I have tried to follow Temkin's advice. I cannot find points of agreement between views which assert and views which deny that we have reasons to care about some things. But when I wrote Part Six of *On What Matters*, I misunderstood some of the people whose meta-ethical views I rejected. Two such people are Railton and Gibbard. I now believe that, as Railton and Gibbard have separately suggested and I shall later try to show, the three of us have resolved our main meta-ethical disagreements. We hope that others will reach similar conclusions.

CHAPTER 38 NON-REALIST COGNITIVISM

130 *Meta-Ethics*

We can roughly distinguish several views that are *meta-ethical* in the sense that they are about the meaning and truth of moral claims, and of other normative claims. Non-Cognitivists believe that most people's normative claims are not intended or believed to state truths. Nihilists or Error Theorists believe that, though these claims *are* intended to state truths, these claims are all false, since there are no normative truths.

Of the Cognitivists who believe that there *are* such truths, some are Normative Naturalists. Normative truths, these people believe, are like other truths about the natural world which might be empirically discovered, in the sense that some partly observable things or events might give us evidence for or against our beliefs in these truths.

According to *Analytical* Naturalists, normative concepts and claims can all be defined or restated in non-normative, naturalistic ways. This

view is fairly plausible when applied to some normative concepts and claims. When people say that something tastes good, or that some act would be rational, these people may mean only that they like this taste, or that this act would achieve the agent's aims. But many normative claims—such as most people's moral claims and many people's claims about reasons—cannot be plausibly restated in non-normative ways. When we say that some act is morally right, for example, most of us don't mean that this act has some natural property, such as that of being an act that would minimize suffering, or an act of which most people would approve.

According to *Non-Analytical* Naturalists, though some concepts and claims are in this sense *irreducibly* normative, these concepts refer to natural properties, and these claims, when they are true, state natural facts. On one such view, though the phrase 'morally right' does not mean 'would minimize suffering', the fact that some act would be right might be the same as the fact that this act would minimize suffering.

According to *Non-Naturalists*, some normative claims state irreducibly normative truths. These truths are not natural, empirically discoverable facts, since we could not have empirical evidence either for or against our beliefs in these truths. When we have decisive reasons to do something, for example, this normative truth could not be the same as some causal or psychological fact, such as the fact that this act would achieve one of our aims. And when certain acts are right, or wrong, these moral truths could not be the same as certain natural facts, such as the facts that these acts would, or would not, minimize suffering, or would be acts of which most people would approve, or disapprove.

131 *Ontology*

Non-Naturalist views can differ ontologically by making different claims about what exists and what is real. *Metaphysical* Non-Naturalists believe that, when we make irreducibly normative claims, these claims imply that there exist some ontologically weighty non-natural entities or properties. Naturalists find such claims mysterious or incredible. *Non-Metaphysical* Non-Naturalists make no such claims, since these people deny that irreducibly normative truths have any such ontologically weighty implications. Nagel, Scanlon, I, and others accept one such

view, which I now call *Non-Realist Cognitivism*. On this view, there are some true claims which are not made to be true by the way in which they correctly describe, or correspond to, how things are in some part of reality. Some examples are the claims that state logical, mathematical, and modal truths, and some fundamental normative truths.

It might be objected that, in distinguishing these views, I have not explained what I mean by ‘ontologically weighty’ or ‘some part of reality’. But I use these phrases when describing views that I don’t accept, and one of my objections to these views is the obscurity of their ontological claims.

We might use the words ‘real’ and ‘reality’ in wider senses, which imply that all truths are truths about reality. If that is how we use these words, it would be less helpful to say that some true claims are not made to be true by the way in which they correspond to reality. But we could restate Non-Realist Cognitivism in a different way. We could say that, on this view, some non-empirical claims do not raise any difficult ontological questions. Mathematicians, for example, should not fear that arithmetical claims might all be false, because there aren’t any numbers.

Gibbard and Blackburn defend a view which they call *Quasi-Realist Expressivism*. The best version of this view, as Gibbard suggests and I shall later argue, is one form of Non-Realist Cognitivism.

CHAPTER 39 NORMATIVE AND NATURAL TRUTHS

132 *Concepts and Properties*

Before considering these views, I shall roughly describe some of the concepts that I shall use. Some people use the word ‘property’ in a narrow sense, which refers to the features of concrete objects or events which can have causes or effects. Two such properties are heat and mass. I use the word ‘property’ in a wider sense. Any claim about something can be restated as a claim about this thing’s properties. Instead of saying that the Sun is bright, or that some argument is valid, we can say that the Sun has the property of being bright, or that this argument has the property of being valid. Since this use of the word ‘property’ adds nothing to the content of our claims, such properties are sometimes called *pleonastic*. Referring to such properties can help us to explain the meaning of some

claims, and to draw some important distinctions. These properties are also *description-fitting* in the sense that they fit the descriptive words or phrases with which we refer to them. Because the word 'luminous', for example, means 'radiates light', the phrase 'being luminous' describes, and thereby refers to, the property of radiating light.

Different descriptive words or phrases, and the concepts they express, may refer to the same property, by describing this property in different ways. Two such concepts are those of *heat* and of *molecular kinetic energy*. In its relevant pre-scientific sense the word 'heat' means, roughly, 'the property that can have certain effects, such as causing certain sensations, melting solids, turning liquids into gases, etc.'. The property that can have these effects, scientists have discovered, is the property of having molecules that move energetically. As this example shows, it can be significant to learn that different concepts refer to the same property in this description-fitting sense.

When different concepts necessarily apply to all and only the same things, these concepts refer to the same property in a different, *necessarily co-extensional* sense. Two such concepts are those expressed by the phrases 'being the only even prime number', and 'being the positive square root of 4'. Since these concepts both necessarily apply only to the number 2, these concepts refer to the same property in this co-extensional sense. But these concepts refer to properties that are different in the description-fitting sense. The phrase 'being the only even prime number' does not describe, and thereby refer to, the property of being the positive square root of 4.

133 *The Co-Extensiveness Argument*

In considering arguments for and against Normative Naturalism, we can take as our example one of the simplest moral views. According to Hedonistic Act Utilitarianism, or *HAU*, acts are morally right if and only if, or *just when*, they minimize the net sum of suffering minus happiness. We need not ask whether this view is true, since most of the claims and arguments that we shall be considering could be restated so that they apply to other views.

Some people defend Naturalism by appealing to the necessity of some normative truths. If *HAU* were true, the concepts *right* and *minimizes*

suffering would necessarily apply to all and only the same acts. These two concepts, some Naturalists argue, would then refer to the same property, which would be the natural property of being an act that minimizes suffering. But this argument does not support Naturalism. Non-Naturalists could reply that, even if these concepts referred to the same property in this necessarily co-extensional sense, these concepts would refer to properties that were different in the more important, because more informative, description-fitting sense.

134 *The Normativity Objection*

According to the *Normativity Objection*, irreducibly normative, reason-implying claims could not, if they were true, state normative facts that were also natural facts. These two kinds of fact are in different, non-overlapping categories. There are many such categories, such as those of physical, logical, legal, musical, and grammatical facts. Just as no fact could be in two of these other categories, it could not be a natural, empirically discoverable fact that we have normative reasons to act in certain ways, or that certain acts are morally wrong. These normative truths could not be explained in naturalistic ways, nor could there be any empirical evidence either for or against our beliefs in these truths.

The Normativity Objection need not assume that, as Non-Naturalists believe, there are some irreducibly normative truths. Many Non-Cognitivists and Error Theorists also believe that some normative concepts and claims are in a separate, distinctive category, so that these claims could not state natural facts. These people add that, since all facts are natural, there are no normative facts.

135 *Scientific Analogies*

When some Naturalists reply to the Normativity Objection, these people appeal to cases in which words with quite different meanings, and the concepts they express, refer to the same property. These Naturalists often give, as their examples, the facts that water is H_2O , and that heat is molecular kinetic energy. These analogies do not, I argue, support Naturalism. These cases show only that, when different concepts correctly describe and thereby refer to the same property, this fact may not

be directly implied by these concepts, so that we may have to discover this fact, or come to know this fact in some other way. Naturalists cannot reject the Normativity Objection by claiming that truths about the identity of properties may not depend on the concepts with which we refer to these properties. That claim is false. To answer the Normativity Objection, Naturalists would need first to explain how it might be true that some irreducibly normative concepts refer to natural properties. They would then need to show how we use these normative concepts to make irreducibly normative claims which, if they were true, would state natural facts.

The scientific analogies are in one way helpful here. In the relevant, pre-scientific sense, 'heat' means, roughly, 'the property, whichever it is, that has certain effects'. Scientists discovered that this property is molecular kinetic energy. There is a similar normative concept which we can express with the phrase: *the natural property, whichever it is, that makes acts right*. If HAU were true, this normative concept would refer to the natural property of minimizing suffering. But this fact, I argued, would not support Naturalism. If this normative concept referred to this natural property, that would not imply that being an act that minimizes suffering is the same as being right. We should instead claim that, if HAU were true, being an act that minimizes suffering would be the natural property that made acts have the different, normative property of being right.

I earlier suggested how Naturalists might reject these claims. If there are certain natural properties that would *make* acts right, having these properties would not *cause* these acts to be right. It is similarly true that, if some object has molecular kinetic energy, that does not cause this object to be hot, and that if some liquid is composed of H_2O , that does not cause this liquid to be water. These Naturalists might claim that, as these cases show, the relation of *non-causal making* implies *being the same as*. When some object has molecular kinetic energy, this fact both makes this object hot and is the same as this object's being hot. When some liquid is composed of H_2O , this fact both makes this liquid water and is the same as this liquid's being water. It is similarly true, Naturalists might claim, that if there is some natural property which is the property that makes acts right, this natural property would be the same as the

property of being right. When I rejected this argument in my *Volume Two*, I did not explain clearly enough how these analogies fail to support Naturalism. I try to do that here.

136 *The Triviality Objection*

According to Non-Analytical Naturalists, though we make some irreducibly normative claims, these claims, when they are true, state natural facts. Such views take two forms. *Hard Naturalists* believe that, since all facts are natural, we don't need to make any such irreducibly normative claims. According to *Soft Naturalists*, we do need to make such claims. Though true normative claims could state only natural facts, having true normative beliefs about these facts would help us to make good decisions and to act well.

Soft Naturalism, I argued, could not be true. If there were no irreducibly normative truths, our normative beliefs could not help us to make good decisions and to act well. I called this argument the *Triviality Objection*, but this name is in one way misleading, since it would not be trivial if there were no irreducibly normative truths.

CHAPTER 40 GIBBARD'S OFFER TO NON-NATURALISTS

137 *The Single Property Illusion*

Gibbard claimed that, if it were true that we ought to do something just when this act would maximize net pleasure, the concepts *ought* and *would maximize net pleasure* would refer to the same property. If these properties were one and the same, that would both tell us what we ought to do and explain why we ought to do these things. I assumed that, when Gibbard made these claims, he was defending a version of Soft Naturalism. This defence, I argued, fails.

138 *Naturalistic States of Affairs and Normative Truths*

My remarks were mistaken, since I misunderstood Gibbard's claims. Gibbard was using the phrase 'the same property', not in the description-fitting sense, but in some version of the necessarily co-extensional sense. Gibbard distinguished between the claims that (1) water is water and that (2) water is H₂O. In Gibbard's terminology, these claims signify

the same state of affairs, but they state two different thoughts. Though (1) states a trivial thought, (2) states a significant discovery. Consider next the claims that (3) maximizing net pleasure is the same as maximizing net pleasure, and that (4) maximizing net pleasure is what we ought to do. Gibbard similarly claimed that, though (3) would be trivial, (4) would be, if true, significant. If we use the word 'fact' to refer to the content of a true thought, the true thought stated by (4) would be a normative fact that was distinct from all natural facts.

When Gibbard made these claims, he was not, as I earlier assumed, defending Normative Naturalism. On the contrary, as Gibbard later remarked, he was offering to help Non-Naturalists by suggesting how these people could explain and defend their view. Gibbard pointed out that, even if normative concepts referred in the co-extensional sense to natural properties, we could use these concepts to think irreducibly normative thoughts, which might be about irreducibly normative truths.

CHAPTER 41 RAILTON'S DEFENCE OF SOFT NATURALISM

139 *The Identity of Properties*

Naturalists, I argued, cannot defend their view by appealing to scientific analogies, such as the fact that heat is molecular kinetic energy. In some passages, Railton partly misdescribes my argument. Railton takes me to assume that, if some claim tells us that two different concepts refer to the same property, this claim would be trivial. That is not my view. It was important, I claimed, to discover that the concepts of *heat* and *molecular kinetic energy* refer to the same property, since this discovery told us how this property is related to certain other properties. There are other more particular ways in which these scientific analogies do not support Naturalism.

140 *Railton's First Response to the Triviality Objection*

Soft Naturalists, I argued, cannot defend their claim that, though true normative beliefs would be about natural facts, these beliefs would help us to make good decisions and to act well. Railton suggests that, to answer this objection, these Naturalists might appeal to the complex roles or job descriptions that certain natural properties might fulfil. This reply to the Triviality Objection does not, I claim, succeed. But Railton later responds to this objection in a quite different way.

CHAPTER 42 RAILTON'S RESOLUTION OF OUR DISAGREEMENTS

141 *Railton's Wider View*

As a Metaphysical Naturalist, Railton believes that there are no ontologically weighty non-natural normative properties and truths. But Naturalists, he claims, could believe that there are some *non-ontological* normative properties and truths. Some examples are truths about which acts are wrong, and about which facts give us normative reasons. We could justifiably believe that there are such normative truths, since this belief would not add anything mysterious to our ontology.

142 *What is Achieved by Railton's Wider View*

In adding these claims to his view, Railton has not given up any of his earlier positive beliefs. These new claims do not conflict, for example, with Railton's earlier response-dependent accounts of some moral, prudential, and aesthetic truths or values, or with his claims about some rule-involving normative truths. Railton has merely come to believe that there are some normative truths of a different kind that he, and many other philosophers, had not earlier considered.

Railton's wider view avoids or answers all of my objections to Normative Naturalism, such as the Normativity and Triviality Objections, and what I called the Soft Naturalist's Dilemma.

Railton's view has also become simpler and more straightforward. Railton earlier claimed that if the normative concept *morally right* referred to some natural property, such as that of minimizing suffering, this property would have to do *double duty*, by being both descriptive/explanatory and normative. It is hard to see how some natural property could do such double duty. Railton now distinguishes between certain normative truths and the natural truths by which these normative truths are non-causally made to be true. These claims are easier to understand and to defend. There are other ways in which, by appealing to this wider view, Railton better achieves his philosophical aims. He can defend claims about what matters, not in a merely response-dependent sense, but in the stronger sense that we all have reasons to care about these things. He can also defend stronger claims about moral truths.

143 *Railton's Commentary*

After reading these remarks of mine, Railton wrote: 'Derek Parfit's response to my paper is, to me, immensely heartening. We are indeed climbing the same mountain.' I shall not try to summarize Railton's Commentary, which I find immensely heartening.

CHAPTER 43 JACKSON'S NON-EMPIRICAL NORMATIVE TRUTHS

144 *Jackson's Co-Extensiveness Argument*

In defending Normative Naturalism, Jackson argues that, since normative concepts are necessarily co-extensive with certain naturalistic concepts, these concepts refer to the same natural properties. Non-Naturalists can reply that, though these concepts would refer to the same properties in the necessarily co-extensional sense, they would refer to different properties in the more informative description-fitting sense. That is how there are some irreducibly normative truths. Jackson's argument does not show that there are no such truths.

When Jackson considers this reply, he concedes that mathematical properties may be able to be individuated in this description-fitting way. But that is not true, he suggests, of the normative properties of people and their acts. This defence of Jackson's argument does not, I claim, succeed.

145 *Jackson's Metaphysical Assumptions*

Jackson assumes that, if there were any non-natural normative truths, these truths would be about ontologically weighty non-natural properties. Given what we have learnt about our world, Jackson claims, we know that there are no such properties and truths. But these normative truths are not, I have claimed, about such ontologically weighty non-natural properties. As some of Jackson's other claims seem to imply, Metaphysical Naturalists can consistently believe that there are some non-empirical truths, such as logical, mathematical, and modal truths, and some fundamental normative truths. These truths do not add anything mysterious to a Naturalist's ontology. Jackson and I could therefore resolve our main meta-ethical disagreements.

CHAPTER 44 SCHROEDER'S CONSERVATIVE REDUCTIVE THESIS

146 *Schroeder's Criticisms of the Triviality Objection*

When Schroeder discusses my Triviality Objection to Soft Naturalism, he calls my argument invalid because he misinterprets one of my premises. This misunderstanding may be my fault, since I should have stated this premise in a less ambiguous way. Schroeder suggests several other objections to my argument, but these do not, I claim, succeed.

Schroeder then discusses the version of my argument that applies to Schroeder's reductive thesis about normative reasons. Schroeder replies that, to answer this argument, he could restate his thesis so that it makes claims about *weighty* reasons. This reply does not, I claim, succeed.

147 *How Schroeder and I Could Resolve our Disagreements*

After reading the remarks that I have just summarized, Schroeder objected that I had misunderstood his view. This objection seems to me justified. After rereading Schroeder's book *Slaves of the Passions*, I now believe that Schroeder's meta-ethical views may not conflict with mine. I failed to realize that Schroeder was not discussing questions about the cognitive significance of our normative beliefs. Schroeder's arguments do not imply that we cannot have true irreducibly normative beliefs.

Schroeder also hopes to defend the view that some things matter, in the sense that we have reasons to care about some things. This view would be hard to defend if, as Schroeder sometimes claims, all of our reasons were given by facts that are in part about our present desires. The fact that we have certain desires could not give us reasons to have them. But Schroeder could revise this part of his view while keeping most of his other claims.

148 *How I Mised Russell*

Russell mistakenly assumes that, on my view, all conceptual truths are trivial. This misunderstanding is my fault, since I failed to repeat my earlier remark—buried in a parenthesis in an endnote—that some conceptual truths are not trivial. I also misled Russell by using the word 'might' in an ambiguous way. I accept many of Russell's other claims.

PART EIGHT EXPRESSIVIST TRUTHS

CHAPTER 45 QUASI-REALIST EXPRESSIVISM

149 *Desires, Attitudes, and Beliefs*

To explain the meaning of our normative claims, Expressivists believe, we should describe the states of mind that these claims express. In claiming that some act is wrong, for example, we are expressing an attitude of being against such acts. Expressivists are *Quasi-Realists* if they add that such normative claims can be true. When we call some claim true, some Minimalists argue, we are merely repeating this claim. For Quasi-Realism to be a distinctive meta-ethical view, Quasi-Realists must use the word 'true' in some stronger, more-than-minimal sense. I briefly describe one such sense.

I earlier assumed that, on Blackburn's view, when we make some claims which seem to be meta-ethical, we are really making first-order, normative claims. That is not, Blackburn claims, his view. When Expressivists discuss meta-ethical questions, Blackburn writes, they could say that being good is what it is to be good, and being wrong is what it is to be wrong. Though such answers are boring, they 'ought to be enough'. These answers are not, I claim, enough. Blackburn also claims that, if two normative judgments express desires or goals that are incompatible, in the sense that they cannot both be fulfilled or achieved, one of these judgments must be mistaken. This claim, I argue, is not true.

150 *Earning the Right to Talk of Moral Truth*

When we believe that some act is wrong, most of us assume that our belief is, or at least might be, true. If Expressivists deny that such beliefs might be true, they should become Error Theorists. Quasi-Realist Expressivists could instead widen their view. They could claim that, when we say that some act is wrong, we both express an attitude of being against such acts, and claim that, in having this attitude, we are getting things right. If we *are* getting things right, such claims would be true. This wider version of Quasi-Realism would be one form of Cognitivism. That would not be, as I mistakenly claimed, an objection to this view. Quasi-Realists could

reply that, in explaining how such expressivist normative beliefs can be true, they had achieved one of their main aims. They would have shown how, in Blackburn's phrase, we can earn the right to talk of moral truth.

Quasi-Realists also believe that true normative claims are not made to be true by the way in which they correctly describe how things are either in the natural world, or in some other non-natural part of reality. This wider version of Quasi-Realism, we can therefore claim, is an Expressivist version of Non-Realist Cognitivism.

CHAPTER 46 GIBBARD'S RESOLUTION OF OUR DISAGREEMENTS

151 *Gibbard's Convergence Claim*

Gibbard himself suggests that his view should take this wider, Non-Realist Cognitivist form. Our normative concepts and claims, Gibbard believes, cannot be defined or restated in naturalistic terms. As Non-Naturalists believe, these concepts and claims are irreducibly normative. According to Metaphysical Non-Naturalists, these claims imply that there exist some ontologically weighty non-natural entities or properties. Gibbard rejects such views, which he finds mysterious and incredible. In his latest book, however, Gibbard makes a striking positive claim. Gibbard suggests that, if Non-Naturalists gave up their ontological beliefs in these mysterious non-natural properties, the best version of Non-Naturalism would coincide with the best version of Gibbard's Quasi-Realist Expressivism.

Some Non-Naturalists do not have such ontological beliefs. According to Non-Realist Cognitivists such as Nagel, Scanlon, and me, there are some non-empirical truths which have no weighty ontological implications. Some examples are logical, mathematical, and modal truths, and some fundamental normative truths. We accept the version of Non-Naturalism that Gibbard suggests would coincide with the best version of Gibbard's view.

152 *Does it Matter Whether Things Matter?*

Some things matter, I claimed, only because there are some non-natural reason-involving normative truths. Gibbard argues that, since it is certain that some things matter, but not certain that there are any such normative

truths, things would matter even if there were no such truths. This argument is not, I claim, valid.

Gibbard also claims that, compared with questions about which are the things that matter, it matters less whether mattering involves non-natural, reason-involving truths. These questions overlap. What matters depends upon the ways in which things matter. It also matters whether these ways of mattering involve such non-natural truths.

After reading an existentialist novel, a young friend of Richard Hare's concluded in despair that nothing matters. It is a difficult question whether and how it matters whether anything matters. If we believe that suffering matters, we may regret this fact. We might try to believe that, as Nihilists claim, nothing matters, because we have no reasons to care about anything. We might then conclude in despair that Nihilism is false, because some things, such as suffering, really do matter.

153 *Getting it Right*

On the simplest version of Gibbard's view, when we say that suffering matters, we are saying: 'Care about suffering!'. Such imperatives could not be true, so this kind of mattering does not involve either natural or non-natural truths. But though imperatives cannot be true, they may get things right. Gibbard himself claims that, in telling people to care about suffering, we are getting it right. If we claim to be getting it right, and our claim is true, suffering would matter in this wider expressivist sense. When he discusses our moral beliefs, Gibbard could similarly claim that, when we say that some act is wrong, we are both expressing the imperative 'No one ever act like that!', and claiming that, in accepting and expressing this imperative, we are getting it right. If we *are* getting it right, these moral claims would be true. To explain how we might be getting things right, Gibbard might appeal to some of his beliefs about normative reasons.

154 *Gibbard's Metaphysical Doubts*

Truths about normative reasons, Gibbard assumes, would not be natural, empirically discoverable facts. Gibbard earlier suggested how we could defend the claim that there are some non-natural normative facts. We could distinguish between states of affairs, which are all naturalistic,

and irreducibly normative thoughts. If we use the word ‘fact’ to refer to the contents of true thoughts, Gibbard wrote, there are perhaps some normative facts that are distinct from all natural facts. Gibbard may have used the word ‘perhaps’ because he is a Metaphysical Naturalist who doubts whether there could be any such non-natural normative facts, or truths. Gibbard also suggests that, if there were such non-natural truths, they would at best be second-rate facts.

In considering these doubts, we can again compare these normative truths with some other non-empirical truths, such as logical, mathematical, and modal truths. These other truths are not second-rate facts. Two plus two must equal four and could not possibly equal three or five. Nor do these truths involve ontologically weighty non-natural entities or properties of the kind that Gibbard finds mysterious. Mathematicians need not fear that, because numbers don’t exist in the spatio-temporal world, and there is no other part of reality in which numbers might exist, arithmetical claims might all be false because there aren’t any numbers.

Similar claims apply to non-empirical normative truths. We need not fear that no acts are wrong, or that we have no reasons to have beliefs or desires, because there are no mysterious non-natural properties of being wrong or being a normative reason. When Gibbard first explained and defended his Expressivist theory, his stated aim was to save what is clear in our normative thinking with one exception, which is our tendency to believe in the existence of mysterious, ontologically weighty non-natural properties. Gibbard also wrote that the appeal of such Platonistic beliefs comes chiefly from a lack of anything to put in their place. Non-Realist Cognitivism is what we can put in their place.

CHAPTER 47 ANOTHER TRIPLE THEORY

155 *Gibbard’s Commentary*

I shall not try to summarize Gibbard’s Commentary.

156 *A Happy Ending*

Nor shall I summarize my response.

PART NINE NORMATIVE AND PSYCHOLOGICAL REASONS

CHAPTER 48 EXPRESSIVIST REASONS

157 *Blackburn's Bafflement*

Blackburn is baffled by my view, which he finds bizarre and ludicrous. I have been baffled by some of Blackburn's claims. I suggest some ways in which we have misunderstood each other.

158 *Blackburn's Beliefs about Reasons*

When we say that someone has a reason to act in some way, Williams claimed, we mean roughly that this act might fulfil one of this person's present desires, or that after informed deliberation this person would be motivated to act in this way. Williams called this the *internal* sense of the phrase 'a reason', and he doubted whether we could intelligibly use this phrase in any other *external*, purely normative sense. Williams imagined a man who is cruel to his wife, and who has no such internal reasons to treat his wife better. Since Williams believes that all reasons are internal, he concludes that this man has no reason to treat his wife better. When Blackburn discusses this example, he rejects Williams's conclusion. We could tell this man, Blackburn claims, that he has decisive external reasons to treat his wife better. In this disagreement, it is the Externalists who win.

This Externalist victory, Blackburn then writes, 'is entirely hollow'. Though we are free to claim that other people have decisive reasons to act in certain ways, even when these acts won't fulfil any of these people's present desires, such claims won't help us to achieve our aim of getting these people to act in these ways. As these and some of Blackburn's other claims suggest, Blackburn and I may not use the same concept of an external, purely normative reason.

CHAPTER 49 SUBJECTIVIST REASONS

159 *Smith's Defence of Subjectivism*

I earlier claimed that, on Subjectivist views about reasons of the kind that Smith has earlier defended, we cannot have any reason to want to

avoid future agony for its own sake, or as an end. If we ask ‘Why not?’, Subjectivists have no good reply. Smith replies that, if we were fully procedurally rational, we would want to avoid future agony because such agony would interfere with our exercise of our rational capacities. This reply does not explain why we can’t have any reason to want to avoid agony, not as a means of fulfilling some other present desire, but as an end, or for its own sake. Smith also argues that, unless the concept of a reason to have some desire can be reduced to the concept of a reason to have some belief, we can’t have any reasons to have desires. Smith’s ingenious argument does not, I claim, succeed. We can plausibly revise Smith’s first premise, and this argument then counts against Smith’s view.

160 *Street’s Defence of Subjectivism*

Street rejects my claim that Williams didn’t understand the concept of a purely normative reason. On a more charitable interpretation, Street suggests, Williams understood this concept, and merely disagreed with me about which facts could give us such reasons. My interpretation, I believe, is more charitable. I accept Williams’s claim that he didn’t understand the concept of what he called an external reason, and I accept all of Williams’s claims about what he calls internal reasons. It is not uncharitable to believe that all these claims are true. Street makes some other claims about reasons, which I question.

CHAPTER 50 STREET’S META-ETHICAL CONSTRUCTIVISM

161 *Street’s Debunking Arguments*

According to Street’s evolutionary debunking argument, because our normative beliefs were greatly influenced by natural selection, these beliefs were caused in ways that were unrelated to their truth. When we know that our beliefs were caused in such ways, these beliefs cannot be justified. Street’s debunking argument, I claim, applies to her own normative beliefs. Street could plausibly deny that her argument undermines these beliefs. We can similarly deny that Street’s argument undermines the most important normative beliefs that Street claims cannot be justified.

162 *Street’s Relativism*

Street argues that, to defend some of her claims, she can accept and defend some relativist view about truth, normativity, and reasons. Street’s

defence of this view is predictably bold and subtle, but it does not, I argue, succeed.

163 *The Normative Implications of Street's View*

Street also claims that, if we accept her meta-ethical view, our deepest evaluative convictions should remain untouched. That is not, I believe, true. Street imagines a monstrous man who believes that, as his coherent attitudes imply, torturing others for his own amusement is the highest value, and is how it is best to live. Street claims that, in having these beliefs, this imagined man couldn't be making a mistake, or be missing something. If this man couldn't be failing to get things right, or be missing something, that would have to be because there is nothing to get right, and nothing to miss. On this view, there couldn't be better or worse ways to live. If we accepted this view, our deepest evaluative convictions would not remain untouched, but would be undermined. Street's ingenious and forceful arguments, though I believe them to be unsound, help us to make philosophical progress.

164 *Chappell's Claims about Street*

When I rejected Street's evolutionary debunking argument, I rejected one of Street's premises. Though our normative beliefs were influenced by natural selection, this influence was not great enough, I claimed, to give strong support to Street's evolutionary debunking argument. Chappell responds to Street's debunking argument in a bolder way, by suggesting that the causal origins of our normative beliefs could not epistemically undermine these beliefs. Chappell also suggests that, in responding to the Skeptic's Argument from Disagreement, we can deny that it makes a difference whether, in ideal conditions, we would *in fact* have similar normative beliefs. I question these suggestions.

CHAPTER 51 MORALITY, BLAME, AND INTERNAL REASONS

165 *Darwall's Claims about Internal Reasons*

No fact could give us a reason, Darwall claims, if we could not possibly be aware of this fact, or our awareness of this fact could not possibly motivate us. Darwall assumes that, since my view about reasons is 'steadfastly externalist', I would reject these claims. That is not so. What I deny is the different claim that no fact could give us a reason unless it is true

that, after informed and procedurally rational deliberation, our awareness of this fact would *actually* motivate us.

166 *Darwall's Defence of Moral Internalism*

The moral wrongness of an act, Darwall claims, cannot give us a reason if we could not possibly believe that such acts are wrong. Those whom Darwall calls *Externalists* about reasons could accept this claim. Darwall also claims that an act's wrongness cannot give us a reason if our belief that this act is wrong could not possibly motivate us. As before, Externalists like me could accept this claim. What we reject is the claim that an act's wrongness cannot give us a reason unless our belief that this act is wrong would *actually* motivate us. That is a very different claim.

On the view that Darwall elsewhere calls Moral Judgment Internalism and finds plausible, we could not believe that some act is wrong without being to some degree motivated not to act in this way. If that were true, we need not consider cases in which such moral beliefs would not actually motivate us, since there could not be any such cases.

Darwall makes several plausible claims about moral accountability and blameworthiness. These claims, I argue, do not conflict with the true beliefs about reasons that Darwall calls Externalist and I call Objectivist.

CHAPTER 52 NIETZSCHE'S MOUNTAIN

167 *Nietzsche and the Convergence Claim*

Huddleston questions my attempts to reconcile some of Nietzsche's claims with what most of us believe. Though some of my attempts may fail, Nietzsche's beliefs do not cast doubt on my Convergence Claim.

PART TEN ETHICS

CHAPTER 53 WHAT MATTERS AND UNIVERSAL REASONS

168 *The Any-All Thesis*

Reasons are *person-neutral* if they are reasons for everyone to have and to try to achieve the same common aims, and *person-relative* if they are

reasons for different people to have and to try to achieve different aims. Reasons are *impartial* if they are reasons that everyone would have, even if their point of view was impartial. *Neutralists* claim that all reasons are person-neutral and impartial. On one such view, we always have most reason to do whatever would be, on balance, best for everyone. *Personalists* claim that all reasons are person-relative. On one such view, we always have most reason to do whatever would be best for ourselves. I believe that, as *Dualists* claim, we have both kinds of reason.

On Nagel's Dualist view, when we have personal reasons to try to achieve certain kinds of aim, other people have corresponding but much weaker impartial reasons to want us to achieve these aims. But some of our reasons, Nagel believes, are purely personal in the sense that no one else has any such corresponding impartial reasons. I believe that, as *the Any-All Thesis* claims, there are no such purely personal reasons. We have no reason to try to achieve some aim if this aim's achievement would not be in any way good. When we have some aim whose achievement would be in some way good, everyone has a weak impartial reason to want us to achieve this aim.

169 *Universalism about What Matters*

There are some plausible counterexamples to the Any-All Thesis, but these are not decisive. If we can defend this thesis, we can be Universalists about what matters. We can believe that we all have reasons to care about the same things. These are the things that matter.

CHAPTER 54 CONFLICTING REASONS

170 *Sidgwick's Problem*

Sidgwick believed both that we always have most reason to do our duty, by doing whatever would be impartially best, and that we always have most reason to do whatever would be best for ourselves. These beliefs imply that, when one act would be impartially best but another act would be best for ourselves, we would have most reason to act in each of these different ways. That is a contradiction, which couldn't be true. Sidgwick's beliefs can be revised so that they avoid this contradiction. But these beliefs would still imply that, whenever any impartial moral

reason conflicts with any self-interested reason, neither reason would be stronger. Reason would give us no guidance, since there would be nothing that we had more reason to do. Sidgwick called this the profoundest problem in ethics.

171 *Moral and Self-Interested Reasons*

De Lazari-Radek and Singer suggest that, to solve Sidgwick's problem, we can give an evolutionary debunking argument against the belief that we have strong self-interested reasons. Most of us care about our own well-being much more than we care about the well-being of strangers. Natural selection explains this fact, since those early humans whose genes made them more self-interested would have been more likely to survive and spread these genes. When our ancestors came to believe that they had most reason to do what would be best for themselves, this belief merely endorsed these self-interested motives. These motives and this belief would have been reproductively advantageous whether or not this belief is true, so we were caused to have this belief in a way that was unrelated to the truth. This fact, de Lazari-Radek and Singer claim, casts strong doubt on this belief. No such argument applies to the belief that we have impartial reasons to do what would be on the whole best for everyone. This belief, and the impartial motives that this belief endorses, would not have been reproductively advantageous. This evolutionary argument against Rational Egoism has some force, but is not decisive.

Many self-interested reasons are decisively outweighed by conflicting impartial reasons, or by other moral reasons. But we often have sufficient self-interested reasons to do what would make things go worse, and we could sometimes have sufficient self-interested reasons to act wrongly.

172 *Other Problems*

Sidgwick believed that impartial reasons never conflict with moral reasons. If these reasons sometimes conflict, as most of us believe, these cases raise some other problems. We might have sufficient or even decisive impartial reasons to act wrongly.

As one example of such conflicting reasons, we can suppose that, as the *Means Principle* claims, it would often be wrong to kill one person as

a means of saving the lives of several other people. On one version of this principle, if some nation is fighting a just war, this nation's government may justifiably use *tactical* bombing against military targets, with the foreseen side effect that some civilians will be killed, but it would be wrong to use *terror* bombing, which kills civilians as a means of persuading the enemy to surrender.

We can apply this principle to an imagined case which is fairly similar to the actual state of the world in 1945. The US President, we can suppose, can choose between two ways of ending the Second World War. In the *Nuclear Policy*, an atomic bomb would be dropped on some Japanese city, killing about 100,000 civilians. This policy would swiftly end the war by persuading the Japanese Government to surrender. In the *Conventional Policy*, the US armed forces would invade Japan, and win the war with much bombing and fighting, whose foreseen side effects would be to kill about 300,000 civilians.

If we accept the Means Principle, we shall believe that the Nuclear Policy would be wrong, because it would involve the killing by pure terror bombing of very many people. But though the President would have a strong moral reason not to act wrongly, by choosing the Nuclear Policy, he would also have a strong conflicting impartial reason to choose this policy, which would be given by the fact that this way of ending the war would kill 200,000 fewer civilians. This impartial reason would not be weaker than, and might be stronger than, the President's moral reason not to choose this policy. The President would have a sufficient and perhaps decisive reason to act wrongly. That would be a disturbing conclusion.

CHAPTER 55 THE RIGHT AND THE GOOD

173 *Moral Ambivalence*

If we are not Act Consequentialists, we may think:

- (1) We could always reasonably want and hope that things will go in the ways that would be best.
- (2) It would often be best if some people acted wrongly.

Therefore

- (3) We could often reasonably want and hope that some people will act wrongly.

There are two ways of rejecting premise (2). Act Consequentialists claim that, when it would be best if people acted in some way, this fact would make this act right. Some other people claim that the badness of any wrong act would prevent this act from making things go best. This second claim is false. Some wrong acts make things go better by preventing several other similar wrong acts. There are other wrong acts, most of us believe, that would make things go best. *Deontic* badness is the badness that acts may have when and because these acts are wrong. All other kinds of badness are *non-deontic*. In considering this Argument for Moral Ambivalence, we should ask whether and how often the deontic badness of wrong acts would outweigh the non-deontic goodness of their effects.

174 *The Badness of Wrongdoing*

Some people argue that there is little or no deontic badness, since our reasons to prevent some murder would not be much stronger than our reasons to prevent some accidental death. This argument is flawed. If we prevent some attempted murder from succeeding, most of the deontic badness is already there.

175 *Moral and Impartial Reasons*

If we are not Act Consequentialists, we may have to admit that we could often reasonably want and hope that some other people will act wrongly. Suppose again that it would be wrong for the US President to choose the Nuclear rather than the Conventional Policy. The deontic badness of this wrong act would be clearly outweighed by the non-deontic goodness of the fact that 200,000 fewer civilians would be killed. The rest of us could therefore reasonably want and hope that the President would act wrongly. This would be another disturbing conclusion.

176 *Wrongness and Reasons*

Some people claim that when they say that some act is wrong, they mean that we have decisive moral reasons not to act in this way. If this is how

we use the word ‘wrong’, we could claim that no one could ever have sufficient reasons to act wrongly. But this claim would be trivial. On this view, our reasons not to act wrongly would always be decisive because, if we didn’t have decisive reasons not to act in some way, this act wouldn’t be wrong. That is like ensuring that, in any battle, we shall be on the winning side because, if our side is about to lose, we change sides. There is another way in which, if we use the word ‘wrong’ in this decisive-moral-reason sense, that would undermine our moral beliefs. We could no longer ask, as a separate question, whether some act would be wrong. Our question would be only whether we had decisive reasons not to act in this way.

CHAPTER 56 DEONTOLOGICAL PRINCIPLES

177 *The Means Principle*

When people claim that it would often be wrong to do what would make things go best, some of them appeal to our negative duties not to act in certain ways. According to the widely accepted *Harm Principle*, our negative duties not to harm people are much stronger than our positive duties to make things go better by saving people from being harmed. One of two duties would be stronger in the *cost-requiring sense* if we would be morally required to bear greater burdens, if that were necessary, to fulfil this duty. One of two duties would be stronger in the *conflict-of-duty sense* if this duty would be stronger than the other when these duties conflict. Our negative duties not to harm people may be much stronger in the cost-requiring sense than our positive duties to save people from being harmed. But these negative duties are not, I believe, much stronger in the conflict-of-duty sense.

In its simplest form, the Harm Principle implies that it would be wrong to save several people’s lives with some act that would also kill one other person. There is one much discussed imagined case, here called *Side Track*, which has led many people to believe that this principle must be revised. Some driverless runaway train is moving down some track where it threatens to kill five people. You are a bystander who could save these five people’s lives by redirecting this train down some other track where it would kill only one person. The Harm Principle

implies that it would be wrong for you to redirect this train, though you would thereby cause this train to kill four fewer people. Most of us would believe that this act would be justified. In a partly similar case, which I call *Bridge*, another runaway train is threatening to kill five people. You could save these people's lives by causing some other person to fall in front of the train, which would then be stopped by hitting and killing this person. Many of us would believe this act to be wrong.

To explain the moral difference between your acts in *Side Track* and *Bridge*, some people appeal to what we can now call the *Means and Side Effect Principle*. It can be right, these people believe, to save several people's lives with some act whose foreseen side effect would be to kill one other person, but it would be wrong to kill one person as a means of saving several other people's lives.

This distinction has been defended in unconvincing ways. Kamm claims that, given our high status as ends-in-ourselves, it would be wrong for us to be killed as a means of saving other people's lives. We could similarly claim that, given our high status as ends-in-ourselves, it would also be wrong for us to be killed as a foreseen side effect of saving other people's lives. This claim seems as plausible as Kamm's. It is wrong to be killed as a means, some other people claim, because we are being sacrificed without our consent to achieve someone else's goal, or we are being treated not as a person but as a mere thing, or we are being denied the right to choose whether we shall be harmed, or we are not being treated as ends-in-ourselves. Similar objections could be applied, with as much force, to acts that would kill us as a foreseen side effect. Several people appeal to Kant's claim that we must not treat people merely as a means. But we treat people merely as a means, in Kant's special sense, when we fail to treat these people as ends-in-themselves. Kant would not have believed that, if we harm certain people not as a means but only as a foreseen side effect, we thereby treat these people as ends-in-themselves.

178 *Harming and Saving from Harm*

Suppose that the US President learns that an asteroid is approaching the Earth, and is on course to hit some large city, where it would kill many people. The President could justifiably order some missile to be fired

that would redirect this asteroid so that it would hit some less populated area, where it would kill fewer people. As Thomson wrote, when there is such an unintended threat to people's lives, we could justifiably bring it about that 'something that will do harm anyway shall be better distributed'. Harm would be better distributed if it came to fewer people. It is no objection to such acts, Thomson claimed, that we would be interfering with how things are actually going.

These claims conflict with the Harm Principle. That principle claims that, compared with our duty to save people's lives, we have a stronger duty not to kill people. This distinction often depends on how things are actually going. If the President redirects the asteroid away from the large city, he would be merely *saving* the lives of the people who live in this city and he would be *killing* some people who live elsewhere. But this distinction depends entirely on the fact that the asteroid is actually moving towards the large city. Since the people in the large city do not have significantly weaker moral claims, the President could justifiably redirect the asteroid so that it would kill fewer people. Such acts would be justified even if they would cause only slightly fewer people to be killed. In another version of *Side Track*, you could justifiably redirect the runaway train so that it would kill one fewer person.

Suppose next that, in *Hand Grenade*, another runaway train is moving towards several people, whom it threatens to kill. You could save these people's lives, but only by throwing a bomb whose explosion would stop the train, but would also kill some other person who is standing nearby. Many people would believe this act to be wrong. But this act would not be condemned by the Means Principle, since you would kill this single person not as a means but only as a foreseen side effect of doing what would save several other people's lives.

When they consider such cases, some people appeal to another principle. According to *the Redirection Principle*, when there is some unintended threat to several people's lives, we could justifiably save these people's lives by redirecting this threat so that it would kill fewer people, but it would be wrong to save these people's lives by starting some new threat that would kill fewer people. This principle permits your act in *Side Track*, but condemns your act in both *Hand Grenade* and *Bridge*.

We ought, I believe, to reject or revise this principle. Suppose that some fire is threatening to kill ten people. On the Redirection Principle, we could justifiably save these ten people's lives by redirecting this fire so that it would kill five other people, but it would be wrong to save these ten people's lives by starting some flood that would put out this fire but would also kill some other single person. This second act would not, I believe, be wrong. This act would save the ten people's lives in a way that would kill four fewer people, and it is morally irrelevant that the single person would be killed by a flood rather than a fire. Firefighters could justifiably start floods if that would cause fewer people to be killed.

CHAPTER 57 ACT CONSEQUENTIALISM AND COMMON SENSE MORALITY

179 *Good and Bad Ways of Treating People*

When Ross argues against all versions of Act Consequentialism, he claims that

(A) it would often be wrong to treat people in certain ways, such as deceiving or coercing them, or breaking our promises to them, even when such acts would make things go better.

Ross was not an Absolutist, since he believed that such acts would not be wrong if their effects would be sufficiently good. We might justifiably break some promise, for example, if that would enable us to help some injured stranger.

In giving the objection stated by (A), Ross overlooks some versions of Act Consequentialism. Some Act Consequentialists believe that

(B) it would often be intrinsically bad to treat people in certain ways, such as deceiving or coercing them, or breaking promises that we have made to them.

Since Ross's argument ignores the possibility that these acts are intrinsically bad, he assumes that

(C) even when these acts would be wrong, because their effects would not be good enough to justify them, these acts might make things go better.

If we believe (B), however, we might reject (C). Like Ross, we may believe that these acts would be wrong unless their effects would be good enough to justify them. But we may also believe that

(D) if the effects of these acts would not be good enough to justify them, these effects would also not be good enough to outweigh the intrinsic badness of these acts. In such cases, these wrong acts would on the whole make things go worse.

If we have these beliefs, we might also agree with Ross about which of these acts would be wrong. The difference would then be only that, unlike Ross, we believe that the wrongness of these acts could be explained in Act Consequentialist terms.

180 *Deontic and Non-Deontic Badness*

The intrinsic badness of these acts, Ross might object, could not make these acts wrong, because it would be the wrongness of these acts that would make them bad. To answer this objection, we can again distinguish between the *deontic* badness that acts may have when and because these acts are wrong, and all other kinds of badness, which are *non-deontic*. We might then claim that these ways of treating people are non-deontically bad, and that, if these acts did not have good enough effects, their non-deontic badness would make them wrong, so that they would also be deontically bad. These are different kinds of badness, as is shown by cases in which such acts are not wrong, because their non-deontic badness is outweighed by the goodness of their effects. Since these acts would have this intrinsic badness, though they would not be wrong, it couldn't be their wrongness that made them intrinsically bad.

We could next point out that, when we claim that certain acts are intrinsically bad, these claims are similar to Ross's claim that certain acts are *prima facie* wrong. There is a genuine convergence between this version of Act Consequentialism and the Common Sense Morality that Ross defends.

181 *Personal and Shared Duties*

Act Consequentialists, Ross claims, ignore the highly *personal* character of duty. Some of our duties are *person-relative*, or *agent-relative*, in the

sense that different people have duties to try to achieve different aims. We each have duties, for example, to keep our own promises. Some of our duties are also *patient-relative*, in the sense that we have these duties only to certain people, such as those to whom we have made promises, or the people to whom we are related in certain other ways, such as our children, parents, pupils, clients, or those who have benefited us.

We also have some duties that are not in this sense patient-relative, because we have these duties to everyone. Some examples are duties not to kill or harm people. Such duties are often claimed to be person-relative in the first way, by giving different people different aims. We are claimed to have a duty not to kill or harm people, which is different from a duty to cause fewer people to be killed or harmed.

When people make these claims, they are usually thinking of cases in which, if we kill some people, we would cause fewer people to be *wrongly* killed. Such cases are best considered separately, after we have reached some view about what we ought to do in cases in which there are no wrongdoers, since everyone will try to do what they ought to do.

We can first consider our positive duties towards strangers. These duties do not give different people different aims. Rather than having a person-relative duty to save people's lives, we have a person-neutral duty to act in such a way that more people's lives will be saved. It would be wrong for us to save one stranger's life rather than enabling someone else to save two or more strangers' lives.

Similar claims apply to our negative duties. Rather than having a person-relative duty not to kill people, we have a person-neutral duty to act in such a way that the fewest people will be justifiably killed. We could justifiably kill one person when and because we know that, if we don't kill this person, it would become someone else's duty to kill two or more people, which this person would then do. Similar claims apply to other negative duties. We could justifiably deceive or coerce one person if we knew that, if we don't deceive or coerce this person, it would become someone else's duty to deceive or coerce more people.

Turn now to cases that involve wrongdoers. Suppose that some wrongdoer knows that we believe that we have a person-neutral duty to do what would cause the fewest people to be killed. This wrongdoer might then credibly threaten that, if we don't kill some innocent person,

he will kill more people. In such cases, our belief that we had this person-neutral duty might have bad effects. But that doesn't by itself show that we have only a person-relative duty not to kill. When moral theorists claim that these negative duties are person-relative, they often appeal to cases in which, by killing someone, we might cause fewer people to be wrongly killed. But these cases raise distinctive questions. We can plausibly believe, for example, both that we ought to do what would cause the fewest people to be killed, and that we ought not to give in to threats by wrongdoers, since that would encourage future threats in ways that may cause more people to be killed.

We are asking a simpler, wider question. When moral theorists claim that our negative duties are person-relative, they intend this claim to apply to all cases. These people claim that we have a duty not to kill or harm people, which is different from a duty to cause fewer people to be killed or harmed. If our negative duties were in this way person-relative, that would be an important structural difference between these moral beliefs and Act Consequentialism. I have claimed that, in cases that don't involve wrongdoers, these duties are not person-relative. We can plausibly believe that, in such cases, we ought to have the common aims that as many people as possible will be saved from death or other harms, and that as few people as possible will be justifiably deceived, coerced, harmed, or killed.

These remarks do not apply to all our duties. We have some duties which are doubly personal, since these duties give different people different aims, and we have these duties to certain other particular people. But if most of us would believe that our negative duties give us the common aims that I have just described, there would again be less disagreement than there is often assumed to be between Common Sense Morality and Act Consequentialism.

CHAPTER 58 TOWARDS A UNIFIED THEORY

182 *Act Consequentialism*

Some Act Consequentialists assume that it would be best if we all accepted AC and always tried to do whatever would make things go best. But it seems likely that, as Sidgwick believed, things would on the whole go better if most people were not Act Consequentialists but

accepted some improved version of Common Sense Morality. AC might even imply that we ought to try to bring it about that no one accepts AC. This fact would not by itself show that AC is false, but it might indirectly support some other view.

183 *Rule and Motive Consequentialism*

Something is *optimific* if it makes things go best. According to *Rule Consequentialism*, instead of acting in optimific ways, we ought to follow optimific rules. When Rule Consequentialists ask which rules are optimific, some of them consider only what would happen if we *successfully followed* these rules. But on the views that I shall discuss, we ask what would happen if we *accepted* and *tried to follow* certain rules.

How well things go does not depend only on what people do. Even if we always acted in optimific ways, the good effects of our acts might be outweighed by the bad effects of our desires, dispositions, and other such *motives*. According to *Motive Consequentialism*, we ought to have motives that are optimific in the sense that there are no other possible motives whose being had by us would make things go better.

These views can be combined. According to *Rule and Motive Consequentialism*, or

RMC: We ought to have optimific motives and we ought to accept and try to follow optimific rules.

Some people believe that RMC could also be combined with AC. On this view, the optimific motive is that of always wanting and trying to act in optimific ways, as the optimific rule requires us to do. But this view is not, I believe, true.

184 *Optimific Motives and Rules*

Some versions of Rule Consequentialism conflict deeply with Act Consequentialism. Some optimific rules may require us, or permit us, *not* to act in optimific ways. If most of us were pure Act Consequentialists, who were most strongly motivated to do whatever would make things go best, our acts would have many good effects. But our lives would on the whole go better if most of us have some other strong motives and try

to follow some other rules. It is good, for example, that most of us strongly love our close relatives and some friends. Having such love and being loved are some of the greatest goods in most people's lives. The optimific rules would often require us or permit us to act on such optimific motives even when such acts would not make things go best. Similar claims apply to several other widely held moral beliefs. These beliefs can be plausibly defended in Rule Consequentialist ways.

185 *Small Effects and Great Harms*

When we ask whether some act's effects would make this act right or wrong, many of us make serious mistakes. One mistake is the belief that we can ignore very small benefits or harms. Many of us, for example, would believe that

(J) we ought to give to a single person one more year of life rather than giving to each of many people only one more minute of life.

But in some cases (J) would be false. A year is about half a million minutes. If we gave to each of a million people one more minute of life, we would give these people two more years of life. This case is unlikely to occur, but there are many actual cases that are relevantly similar.

Consider next the claim that

(L) most pain could become worse in some way that would be, not merely very small, but imperceptible. In such cases, we couldn't even notice that our pain has become worse.

This claim may seem obviously false. Since pain is bad only because of the way it feels, we may assume that, if our pain doesn't *seem* worse, it can't *be* worse. But that is not so. We can easily show that (L) is true.

Suppose that, in

the Bad Old Days, a thousand torturers each have one victim and one pain-producing machine. At the start of each day, each victim is already feeling mild pain. Each of the torturers turns some switch a thousand times on his machine. Each turning of this switch makes some victim's pain only

imperceptibly worse. But after a thousand turnings each victim is in severe pain, which continues for the rest of the day.

Suppose next that these torturers have moral doubts about what they are doing. One of them suggests that, to answer these doubts, they should connect their machines in a certain way. In the resulting case, which I have called

the Harmless Torturers, each of the thousand torturers pushes some button which turns the switch once on each of the thousand machines. Since all of the switches are again turned a thousand times, the victims suffer the same severe pain. But since each torturer's act turns each switch only once, none of these acts makes any victim's pain perceptibly worse.

These torturers might argue:

It is not wrong to affect someone's pain in some way that is imperceptible.

None of us makes anyone's pain perceptibly worse.

Therefore

None of us is acting wrongly.

This argument's conclusion is clearly false. These torturers are still acting wrongly, since their acts inflict on their victims just as much pain as they inflicted in the *Bad Old Days*. But these torturers can truly claim that none of their acts makes anyone's pain perceptibly worse. To reject this argument's conclusion, we must therefore reject this argument's first premise. We must claim that

(M) it can be wrong to impose pain on people, even if these acts make no one's pain perceptibly worse.

In defending this claim, we might appeal to the effects of each particular act. In some cases we must appeal to such effects. Though these claims are defensible, I shall not discuss them further here.