

The Emotional Construction of Morals

JESSE J. PRINZ

OXFORD
UNIVERSITY PRESS

Preface

David Hume's *Treatise of Human Nature* is divided into three books: "Of the Understanding", "Of the Passions", and "Of Morals". One might wonder how these disparate topics are related, other than by virtue of the fact that they have something to do with the mind. But the links become clear on reading the text. Hume develops a theory of concepts (or "ideas") in the first book and a theory of emotions in the second book, and then he integrates these in the third by arguing that our moral concepts have an emotional foundation. The project is also unified by Hume's allegiance to empiricism. His theory of concepts is based on the premise that ideas are stored copies of sensory impressions, and his theory of emotions is designed to be compatible with this empiricist view (he defines emotions as impressions of impressions). Hume's moral theory is empiricist too. Moral concepts seem especially problematic for an empiricist because there can be no image of virtue, no taste of goodness, and no smell of evil. By appealing to sentiments, Hume is able to argue that all concepts bottom out in impressions, after all. The concept of goodness consists in a feeling of approbation and the concept of badness consists in a feeling of disapprobation. The class of virtues has no common appearance, but good things just feel right; the class of vices would be impossible to paint, but each instance elicits a palpable pang of blame. In sum, Hume's *Treatise* has a coherent structure, and the culminating moral theory can be read as the resolution of an apparent counter-example to his theory of concepts, or as the payoff for those who take the time to understand how the mind works. No matter where you place the emphasis, Hume's theory of concepts and his theory of morals hang together, and passions are the glue.

Philosophers like to reinvent wheels, and I am no exception. The views that I defend here owe a tremendous debt to Hume. This book defends a sentimentalist theory of morality that builds on the ideas developed by Hume and some of his contemporaries. I depart from Hume in various ways, but the basic thrust of the theory is Humean, and, in this respect, my proposals are footnotes to Book III of the *Treatise*. And this is not the first Humean footnote I've written. My first book, *Furnishing the Mind*, defends an empiricist theory of concepts, and my second book, *Gut Reactions*, defends an empiricist theory of emotions (which is more Jamesian than Humean, but, with Hume, my goal there is to show that emotions are a kind of impression). So here, in my third book, I am simply completing a trilogy that parallels the structure of Hume's *Treatise*. These works are independent in one sense—you can reject one while accepting the others—but they hang together in just the way that Hume's *Treatise* hangs together. I view them as parts of a whole, and I view that whole as a tribute and modest extension of Hume's masterwork.

I have three main goals in extending Hume's project. The first is to provide empirical support for a theory that was first developed from an armchair. The second is to add some details to Hume's theory, including an account of the sentiments that undergird our moral judgments, and an account of the ontology that results from taking a sentimentalist view seriously. My third goal is to show that this approach leads to moral relativism. Hume resisted relativism, and I argue that he shouldn't have. I also investigate the origin of our moral sentiments, and I suggest that Nietzsche's genealogical approach to morality has much to contribute here. The resulting story is half Humean and half Nietzschean, but I take the Nietzschean part to fit naturally with the Humean part.

I mention Hume and Nietzsche by way of acknowledgement. Within the pantheon of dead philosophers, they are ones to whom I owe the greatest philosophical debts. I must also mention Edward Westermarck, because he recognized the link between sentimentalism and relativism a hundred years ago, and recognized the value of anthropology and history in investigating morals. This book continues in the tradition of Westermarck. Among living philosophers, I have been especially inspired by Gil Harman, Shaun Nichols, David Wiggins, and John McDowell. Steve Stich also deserves special mention for his efforts to promote an approach to philosophy that makes liberal use of empirical results. On that note, I also owe tremendous debts to the scientists who have been providing data to help assess philosophical theories. Among psychologists, Jon Haidt and James Blair have been an especially influential, and I would also single out the late Marvis Harris, whose cultural materialism leaves its mark on the second half of this book. These authors have educated me through their published work, but many others have offered guidance through discussion and written commentaries on material from this book. I have benefited from giving talks at numerous philosophy departments and conferences, spanning four continents and twice that many countries. I wish I could list the name of everyone who offered suggestions or objections along the way. I also want to thank all the members of the Moral Psychology Research Group, who have created one of the most conducive environments for exchanging philosophical ideas that I have ever seen. I have also benefited from written feedback, which led to improvements large and small throughout. In this context, let me first mention participants in seminars taught by Steve Stich, Eric Schwitzgebel, and John Mikhail who endured earlier versions of this manuscript or related papers. I also received philosophical and typographical corrections on the entire manuscript from Nigel Hope, Mark Jenkins, and Jonathan Prinz, as well as helpful comments on selected parts or related materials from Ruth Chang, Matthew Chrisman, Justin D'Arms, Karen Jones, Matt Smith, Valerie Tiberius, Teemu Toppinen, Brian Weatherson, and others whom I am undoubtedly forgetting. Among readers, my biggest debt goes to Shaun Nichols, Richard Joyce, and two anonymous referees for Oxford University Press, who provided me with detailed comments on drafts of the manuscript. They each caught embarrassing mistakes and pressed me on

dozens of philosophical issues. The book is much better because of them, and it would have been better still had I been more successful in accommodating all of their suggestions. I will remain forever grateful. Of course, I would not have received such helpful feedback were it not for my patient and outstanding editor, Peter Momtchiloff. Peter has been a great source of support at every stage.

In writing this book, I also benefited from several institutions. I was a fellow at the Collegium Budapest and did some writing there. Tamar Gendler was instrumental in orchestrating that visit, and in assembling a wonderful group of summer colleagues. I also owe special thanks to the Center for Advanced Study in the Behavioral Sciences, in Palo Alto. CASBS is a magical place, and I finished this manuscript there. In so doing, I benefited from the abundant intellectual resources and the outstanding staff, who contribute to making it an ideal environment for research. I was able to go to CASBS because of a research leave from my home institution, the University of North Carolina at Chapel Hill. I am grateful to UNC for that, but also and especially to my students and colleagues. There is no better place to work.

Finally, I wanted to mention my family. I feel fortunate to have been raised by two parents with strong moral convictions, and I grew up alongside an older brother with a keen moral sense. My views about right and wrong would be very different without them, and they continue to provide support in many ways. As always, my deepest gratitude goes to Rachel, who was nearby as I wrote almost every page of this book, and she has patiently endured every mood swing that comes along with the writing process. Her support has been essential.

Contents

Preamble: Naturalism and Hume's Law	1
-------------------------------------	---

I. MORALITY AND EMOTION

1. Emotionism	13
2. Emotions: Non-moral and Moral	50
3. Sensibility Saved	87
4. Against Objectivity	138

II. CONSTRUCTING MORALS

5. Dining with Cannibals	173
6. The Genealogy of Morals	215
7. The Limits of Evolutionary Ethics	244
8. Moral Progress	288

<i>References</i>	309
<i>Index</i>	329

Preamble

Naturalism and Hume's Law

Morality is a normative domain. It concerns how the world ought to be, not how it is. The investigation of morality seems to require a methodology that differs from the methods used in the sciences. At least, that seems to be the case if the investigator has normative ambitions. If the investigator wants to proscribe, it is not enough to describe. As Hume taught us, there is no way to derive an ought from an is. More precisely, there is no way to deduce a statement that has prescriptive force (a statement that expresses an unconditional obligation) from statements that are purely descriptive. No facts about how the world is configured entails that you ought to refrain from stealing or killing or blowing up buildings. Hume's Law is appealing because it makes morality seem special; moral truths are unlike the cool truths of science. But, on one reading, Hume's Law is a recipe for moral nihilism. By insulating moral truths from scientific methods, it may imply that morality is supernatural. If so, morality should go the way of spirits and fairies. That is a path I want to resist.

Defenders of Hume's Law acknowledge the viability of certain kinds of descriptive projects in morality. One can describe the moral convictions that obtain in a culture. One can describe the nature of the concepts that people deploy when they make moral judgments. One can say something descriptive about the nature of moral facts and how they relate to other kinds of facts. These questions will be my concern. But, I want to begin by discussing how the descriptive truths about morality bear on the prescriptive. The metaethical theory and moral psychology that I will be defending in the chapters that follow offers a way to cross the is/ought boundary.

I will argue that morality derives from us. The good is that which we regard as good. The obligatory is that which we regard as obligatory. The 'we' here refers to the person making a moral claim and the cultural group with which that individual affiliates. If the good is that which we regard as good, then we can figure out what our obligations are by figuring out what our moral beliefs commit us to. Figuring out what we believe about morality is a descriptive task *par excellence*, and one that can be fruitfully pursued empirically. Thus, normative ethics can be approached as a social science.

This suggestion is difficult to square with the intuition underlying Hume's Law. There is a nagging intuition that no empirically discoverable facts about our beliefs can entail that we ought to behave in a certain way. I do not want to trample on this intuition. Hume's Law is true in one sense, and false in another. That is what I hope to show here. More precisely, I want to show how a thoroughgoing naturalist—one who is repelled by spirits and fairies—can find a place for the normative. I regard Hume as such a naturalist, and I will be defending a view of morality that is deeply indebted to Hume. The view that I favor preserves many of our intuitions about the moral domain, but not all. I reject nihilism, but embrace subjectivism, relativism, and arationalism. Morality is a human construction that issues from our passions. But that does not mean we ought to give it up.

0.1 FOUR KINDS OF NATURALISM

The term 'naturalism' is used in a variety of ways, sometimes with a derogatory intonation, and sometimes as a battle cry. I want to discuss four different species of naturalism, all of which I support. I will not argue for naturalism here. I will just pledge my allegiance.

One kind of naturalism, already suggested by my remarks about fairies and spirits, is best understood in contrast to supernaturalism. It is the view that our world is limited by the postulates and laws of the natural sciences. Nothing can exist that violates these laws, and all entities that exist must, in some sense, be composed of the entities that our best scientific theories require. This is a metaphysical thesis; it concerns the fundamental nature of reality. I will call it metaphysical naturalism.

Metaphysical naturalism entails a kind of explanatory naturalism. If everything that exists is composed of natural stuff and constrained by natural law, then everything that is not described in the language of a natural science must ultimately be describable in such terms. This is not equivalent to reductionism in the strong sense of that word. Strong reductionists say that the relation between natural sciences and 'higher-level' domains is deductive. We should be able to deduce higher-level facts from their lower-level substrates. Antireductionists deny this. They think, for example, that there are higher-level laws or generalizations that could be implemented in an open-ended range of ways. Regularities captured at a low level would miss out on generalizations of that kind. The explanatory naturalist can be an antireductionist. The explanatory naturalist does not need to claim that low-level explanations are the only explanations. The key idea is that there must be some kind of systematic correspondence between levels. One must be able to map any entity at a high level onto entities at a lower level, and one must be able to explain the instantiation of any high-level generalization by appeal to lower-level features that realize those generalizations.

A third kind of naturalism can be termed methodological. If all facts are, in some sense, natural facts (according to metaphysical naturalism), then the methods by which we investigate facts must be suitable to the investigation of natural facts. Philosophers sometimes claim to have a distinctive method for making discoveries: the method of conceptual analysis. If metaphysical naturalism is true, this cannot be a supernatural method of discovering supernatural truths. Concepts themselves are natural entities, and they can be investigated using natural processes. Conceptual analysis is, like all legitimate investigatory tools, an empirical method. As empirical methods go, it is not especially powerful. Conceptual analysis proceeds through first-person access to psychological structures, or introspection. Introspection is error-prone, and there are methodological perils associated with drawing conclusions from investigation using a single subject (oneself). We can investigate concepts using the tools of social science. If concepts are natural entities, then they come about in natural ways. For example, concepts can be acquired through experience, and they can be revised through experience. They have no special status when it comes to revealing facts about the world.

Methodological naturalism, as I have defined it, is associated with Quine. In his (1969) critique of epistemology, Quine tells us that the investigation of knowledge should be pursued using the resources of the social sciences. In his (1953) defense of confirmation holism, Quine argues that all claims are subject to empirical revision. There is a further kind of naturalism associated with Quine's holism. We are always operating from within our current theories of the world. In making theoretical revisions, we cannot step outside our theories and adopt a transcendental stance. To do so would be to suppose that we have a way of thinking about the world that is independent of our theories of the world. If theories of the world encompass all of our beliefs, then no such stance is possible. Call this transformation naturalism, because it is a view about how we change our views.

Each form of naturalism has implications for normativity. Metaphysical naturalism entails that moral norms, if they exist, do not require postulating anything that goes beyond what the natural sciences allow. Explanatory naturalism entails that we can ultimately describe how any moral norm is realized by natural entities. Methodological naturalism entails that we should investigate norms using all available empirical resources tools. Transformation naturalism entails that we must investigate norms from within our current belief systems, and, as a result, the norms we currently accept will influence our intuitions about what norms we ought to uphold. If we chose to change our norms, we cannot do so by adopting a transcendental stance that brackets off the norms we currently accept.

0.2 BREAKING HUME'S LAW

If naturalism is right, then moral facts are natural facts, or they are not facts at all. Natural facts are facts that are consistent with the four strictures of naturalism

just adduced. The world is as it is, and not any other way. If the world includes facts about what ought to be, those facts must be explicable in terms of how things are. Every ought must supervene on an is. Since naturalism does not entail reductionism, naturalism does not entail that prescriptive facts reduce to descriptive facts. Naturalism does, however, entail that prescriptive facts are descriptive facts in another sense. Every prescriptive fact must be realized by, or made true by, facts that can be described without use of prescriptive vocabulary. For every prescriptive fact there is some underlying descriptive fact that makes it true. As it happens, I think that naturalism does allow us to infer prescriptive facts from normative facts, and, thus, there is a way to break Hume's Law. But naturalism does not entail that Hume's Law is violable, for reasons that I will discuss in the next section.

First, I want to offer a quick and dirty argument for how to derive an ought from an is. A full defense of the argument would require a more labored excursion into the philosophy of language. My goal here is more modest. I want to indicate one way in which a naturalist might simultaneously regard moral facts as natural (hence entailed by descriptive facts), but also irreducible (and thus not so entailed). The arguments in this section and the next illustrate how that seemingly paradoxical pair of demands might be met.

To see how an ought might be derived from an is, we must first figure out what oughts are. The way to do that is to figure out what the word 'ought' means (here I restrict myself to the moral use of 'ought'). What concept does that word express? To answer this question, we need to do some psychology (introspective or otherwise). We need to determine what people have in mind when they say that something is obligatory. Much of this book is about that question. For now, I want to sketch a very simplified version of the kind of answer that I will defend. On the theory I favor, when a person says that a course of action is obligatory, that judgment expresses what might be called a prescriptive sentiment. A prescriptive sentiment is a complex emotional disposition. If one has this sentiment about a particular form of conduct, then one is disposed to engage in that conduct, and one is disposed to feel badly if one doesn't. One is also disposed to condemn those who don't engage in that form of conduct. Suppose that Smith honestly judges that one ought to give to charity. Smith is expressing a sentiment that disposes him to feel badly if he doesn't give to charity and angry if you don't give to charity. This resembles the philosophical view called emotivism, but, as will become clear in chapter 3, my approach differs in important details.

Many refinements will follow in the coming chapters. I want to dwell here on implications. If the word 'ought' expresses a prescriptive sentiment, then that is what the word means. The concept underlying the word can be nothing more than what we use the word to express. So, if this simplified psychological theory is right, then we have learned what it means to say that someone ought to do something. We have learned what conditions satisfy the judgment that something is obligatory.

Now we are in a position to try to get an ought from an is. I offer the following argument:

1. Smith has an obligation to give to charity if ‘Smith ought to give to charity’ is true.
2. ‘Smith ought to give to charity’ is true, if the word ‘ought’ expresses a concept that applies to Smith’s relationship to giving to charity.
3. The word ‘ought’ expresses a prescriptive sentiment.
4. Smith has a prescriptive sentiment towards giving to charity.
5. Thus, the sentence ‘Smith ought to give to charity’ is true.
6. Thus, Smith has an obligation to give to charity.

The conclusion of this argument is a prescriptive fact. The premises are descriptive. The word ‘ought’ is mentioned, but never used. Hume’s Law has been violated.

My argument contrasts with an argument defended by Searle (1964). Searle also pursues a metalinguistic strategy. Simplifying a bit, he says that, when a person utters a sentence of the form, ‘I promise to do X’, that person places herself under an obligation. This is part of the meaning of promising. Then Searle infers that a person who has placed herself under an obligation is under that obligation. I am not convinced by Searle’s argument. There may be trouble with both steps (for a more thorough critique, see, e.g., Downing, 1972). To promise is only to place oneself under an obligation if people ought to keep their promises. Thus, there is a suppressed normative premise. The move from placing oneself under an obligation to being under an obligation is also suspect. Placing oneself under an obligation can be interpreted conventionally. It can be a matter of being regarded as falling under an obligation in the eyes of a community. The community can regard a person as having an obligation—can place her under an obligation—even if the person is not actually obligated.

I think we need a stronger metalinguistic premise than Searle offers. We need a substantive theory of the meaning of normative terms. Premise 3 in my argument articulates such a theory. That’s where all the action is. The other premises are hard to deny. Premise 3 is controversial, and one goal of the chapters ahead is to provide arguments that make it more convincing. But I hasten to note that the argument can be modified to accommodate other theories. If naturalism is true then moral concepts are either vacuous, or they express properties that can ultimately be described without moral vocabulary. If my analysis of ought is incorrect, substitute another analysis, and replace premise 3 with the corresponding description of the natural facts underlying obligation. Now revise premise 4 accordingly, and the argument will go through. If there are obligations, then they can be derived in this purely descriptive way on any naturalist account.

0.3 SAVING HUME'S LAW

This is all a bit unsettling. First of all, there is an intuition favoring Hume's Law. There seems to be a logical leap from premises about how things are to conclusions about how things ought to be. Second of all, the theory of norms given in premise 3 makes it too easy to derive obligations. A sadistic person might have a prescriptive sentiment towards making people suffer. The argument just presented would entail that the sadist is obligated to be cruel. Something must have gone wrong.

I think these concerns can be addressed. With regard to the first concern, I begin by noting that the argument that I have offered does not violate Hume's Law. The argument does show how we can use descriptive premises to derive prescriptive facts, but the phrase 'prescriptive fact' turns out to be ambiguous. On one reading, a prescriptive fact is just a fact about what someone is obligated to do. But, a prescriptive fact can also be interpreted as a prescriptive judgment or, more succinctly, a prescription. Notice how the conclusion is expressed in the argument above. I said, 'Smith has an obligation to give to charity.' I did not say, 'Smith ought to give to charity.' Indeed, the argument itself shows why this conclusion could not follow. 'Ought' expresses a prescriptive sentiment. It can only be used truly by a speaker who has that sentiment. No premise in the argument entails that I, the author of the argument, have any disposition to react emotionally to charity. So no premise in the argument could entail, in my voice, that Smith ought to give to charity. If 'oughts' are prescriptions, then I have not shown how to derive an ought from an is. Premise 3, which gives the meaning of ought, shows why such a derivation won't work. That premise does not abrogate Hume's Law; it is the key to defending it.

In the end of the last section, I said that Premise 3 could be replaced with premises describing other naturalistic theories of normative terms. Other theories do not necessarily entail the result that I have just presented. They do not necessarily explain why there is no direct inference from obligation to ought. It is an advantage of the approach that I favor that it explains why Hume's Law is so compelling. Normative claims seem as if they can't be derived from descriptive claims, because there is no way to derive a prescriptive sentiment. Identifying normative concepts with prescriptive sentiments captures the truth in Hume's Law.

One might object that my attempt to save Hume cannot work because it violates a basic semantic principle. In the argument above, the final step moves from the semantic premise that 'Smith ought to give to charity' is true, to the claim that Smith has an obligation to give to charity. One might think that the semantic premise entails something stronger. If 'Smith ought to give to charity is true', then Smith ought to give to charity. This is just an instance of disquotation.

We can always infer P from ‘P’ is true. Or can we? I think that the argument that I have presented is a counterexample to the principle of disquotation. This is not a bad bullet to bite, because there are other counterexamples. Suppose Smith utters the sentence, ‘I am Smith.’ That sentence is true. It does not follow that I am Smith. Disquotation is not always allowed when we use indexicals such as ‘I.’ I believe that ‘ought’ is like an indexical in that its meaning is not exhausted by its contribution to a proposition expressed. I will argue for this conclusion in chapter 5. For now, the case of ‘I’ simply shows that disquotation has well-known exceptions. If ‘ought’ is an exception, and if it works like ‘I’, then my argument is sound.

The fact that we cannot derive oughts may come as cold comfort to some. Isn’t it bad enough that we can infer obligations? Inferring obligations from descriptive premises is a little bit disturbing, but I think we can now diagnose why. We are uncomfortable asserting that people have obligations that we do not endorse. We would not want to assert that sadists are obliged to be cruel. I think that this discomfort has a pragmatic origin. Ascriptions of obligations conversationally implicate prescriptive judgments. If I tell you that someone is obligated to give to charity, I probably have an interest in conveying how I feel. Asserting the existence of an obligation is a way of conveying that I think the person ought to do something. But ‘ought’ is a conversational implicature of ‘obligation,’ not a semantic entailment. To see that, notice that the inference from ‘obligation’ to ‘ought’ can be cancelled. It sounds utterly contradictory to say, ‘Smith ought to give to charity, though he ought not to give to charity.’ But it does not sound contradictory to say, ‘Smith has an obligation to give to charity, but he ought not.’ We say things like this quite frequently when talking about the moral values of other people. We might say that the Japanese soldiers of World War II had an obligation to sacrifice their lives as Kamikaze pilots, but they ought not to have done that. Likewise, I can consistently admit that sadists have an obligation to be cruel while insisting that they ought to refrain from cruelty. This addresses the second concern raised at the beginning of this section. Obligations can be deduced from descriptive premises, but they need not be endorsed by their deducers. Endorsements are merely implicated. They cannot be deduced. Believing that Smith ought to give to charity requires making a prescriptive judgment. To make a prescription, we need to be in a particular psychological state—we need to prescribe. That is the sense in which we cannot derive an ought from an is.

0.4 DEFENDING SUBJECTIVISM

I have been arguing that Hume’s Law is basically true. My defense depends on a theory of normative concepts that I presented in the form of a simple sketch. ‘Ought,’ I said, expresses a prescriptive sentiment. My primary goal

in the chapters that follow will be to defend this claim, and to bring out some implications. I will focus on concepts such as GOOD and BAD or RIGHT and WRONG (capital letters denote concepts). These, like the concept OUGHT, essentially involve sentiments. Such concepts are fundamentally *subjective*.

My goal will not be to derive prescriptions from descriptions. That is a normative project and, if the preceding arguments are right, it is not one that can be taken very far. But I will try to derive metaphysical facts from psychological ones. Right and wrong are the referents of our concepts of RIGHT and WRONG if they are anything at all. If the analysis of our concepts uncovers a strong connection to subjective responses, then these terms may refer to something subjective. Moral psychology entails facts about moral ontology, and a sentimental psychology can entail a subjectivist ontology.

If morality is subjective, then why should moral judgments matter to us? One answer, inspired by Hume, is that we can't help caring about morality. There is something right about this, but it only pushes the question back a level. Why can't we help caring about morality? This question may actually be harder to answer than the question of why we *do* care. There is no single answer to the latter question. Moral systems serve various ends. They regulate behavior, they imbue life with a sense of meaning, and they define group membership. The question 'Why does morality matter?' is like the question 'Why does law matter?' or why does 'Culture matter?' People who feel uncomfortable with the idea that morality derives from us, should consider some other things that derive from us, such as medicine, governments, and art. The fact that art is a social construction does not deprive it of value. We don't expect institutions of art to collapse upon discovering that art is a product of human invention.

The discussion ahead divides into two parts, corresponding to themes that emerged in this discussion. In part I, I argue that morality depends on emotions, and, in part II, I discuss what I take to be an implication of this view: the hypothesis that morality varies across cultures. If morality depends on sentiments, I argue, then it is a construction, and, if it is a construction, it can vary across time and space.

The first chapter in part I presents a survey of different ways in which emotions can be involved in morality. I introduce the term 'emotionism' to label any view that makes emotions essential, and I offer some reasons for thinking that a strong form of emotionism is true. In chapter 2, I lay the foundations for an emotionist theory by presenting a general theory of the emotions. If morality has an emotional basis, then it is best to begin with an independently motivated theory of what emotions are. In that chapter, I also present an overview of the moral emotions, and I suggest that moral emotions derive from non-moral emotions. In chapter 3, I begin to present my positive account. It is what contemporary ethicists call a 'sensibility theory,' though my particular version departs in subtle ways from prevailing accounts (namely, it draws on an account of moral sentiments forecast in chapter 2, and it is not metacognitive). I argue

that this theory can cope with ten major objections that have been levied against sensibility theories. Chapter 4 addresses a further objection not addressed in chapter 3: sensibility theories are subjectivist, and many people assume that morality is objective. I argue against this assumption by distinguishing several kinds of objectivity and critically assessing leading ethical theories that purport to show that morality is objective in each sense of the term. I conclude that morality is thoroughly subjective.

I call the account developed in part I ‘constructive sentimentalism.’ The term sentimentalism refers to the role of sentiments, and the term ‘constructive’ refers to the fact that sentiments literally create morals, and moral systems can be created in different ways. Part II focuses on this implication of sentimentalism. More specifically, it explores the role of culture in shaping moral values. In chapter 5, I draw out the relativist consequences of my case against objectivism, and I respond to standard arguments against relativism. The sixth chapter concerns the genealogy of morals, in Nietzsche’s sense. I argue that historical anthropology can be used to explain why certain values persist, and why others have disappeared. I also assess the degree to which such analyses can be used to criticize morality. Chapter 7 turns from genealogy to genes. Even if some values are historical in origin, others may be biological. Evolutionary ethicists have been pushing this line in recent years. I argue that evolutionary ethics falls short of explaining any of our specific values. The only biologically based moral rules are too abstract to guide action, and their status as moral is epigenetic. Morality essentially involves learning. This conclusion bears on the prospect for moral progress, which is the theme in the final chapter. I discuss the nature of moral debates and argue that we can improve on morality. Moral improvement sometimes requires us to look beyond the categories of good and evil, but we should not attempt to abandon morality or replace it with another kind of normative enterprise.

My approach in defending these claims will be naturalistic in all the senses that I characterized above. My most obvious commitment is to methodological naturalism, because I will draw on empirical findings throughout, including findings from neuroscience, psychology, psychiatry, anthropology, cultural history, and ethology. I think enduring philosophical questions can be illuminated by empirical results, and, indeed, they might not endure so long if we use the resources of science. That said, I do not reject traditional philosophical methods, such as conceptual analysis. Indeed, I think that conceptual analysis is an empirical method in some sense: a kind of lexical semantics achieved by means of careful introspection. I think that method often bears fruit, but sometimes introspections clash or fail to reveal the real structure of our concepts. So it is helpful to find other methods to help adjudicate between competing philosophical theories. These other methods cannot replace philosophy. Philosophy poses the problems we investigate, devises useful tools for probing concepts (such as thought experiments), and allows us to move from data to theory by systematizing results into coherent packages that can guide future research. I see

philosophy as continuous with science, and believe that we should be open to using any methods available when asking questions about the nature of morality.

I am also a pluralist about subject matter. This is a book about moral psychology, metaethics, and the origin and anthropology of morals; I even come into contact with some normative questions in the final chapter. Readers with a specific interest in, say, metaethics, may find little of interest in the discussions of cultural history, and readers with an anthropological orientation may be put off by the discussions of moral ontology. I hope this isn't the case. I think a complete account of morality should touch on each of these dimensions, and I think the dimensions are mutually illuminating. For example, one can argue for relativism by presenting semantic evidence and one can argue by studying cultural variation. Both may provide converging evidence, and the cultural observations motivate semantic inquiry and help to reveal why the semantic thesis may be so deeply important.