

Introduction

Global Philosophy and Ethical Theory

Ethical Theory

While we touch on a range of different philosophical traditions in this book, our focus is squarely on what the philosophers from these traditions have to say about “ethical theory” specifically. It is therefore worth considering what is meant by the term “ethical theory.”

Ethical theory is a branch of the broader category of ethical or moral philosophy. In the Western tradition, moral philosophy is traditionally divided into three main branches: applied ethics, normative ethics, and metaethics. Ethical theory is sometimes another name for the second of those branches: normative ethics (though aspects of metaethics can fall under it as well).

Normative ethics can be contrasted with the other two branches of ethics. We can think of these branches as being related to each other in a hierarchy, with metaethics as the most fundamental branch, normative ethics on top of that, and applied ethics at the highest level.

Metaethics deals with the most basic concerns of ethics. What does it mean to call something “moral” in the first place? Is morality something in the world (are there “moral facts”), or is it merely a matter of taste or cultural belief? If there *are* moral facts, what are they, and how do they relate to other things in the world, such as natural facts—facts about the world around us—or human psychology?

At the other end, applied ethics deals in specific moral issues. For instance, if we want to know whether euthanasia is morally permissible, whether it is okay to eat meat, what we owe fellow citizens, or the extent to which we are obliged to act to prevent climate change, we should engage in applied ethics.

Situated between the specific and the basic, we have normative ethics. Normative ethics is more general than applied ethics. It is not directly concerned with specific issues, but rather with how we *work out* what we should do. In this sense, we might think of normative ethics as standing behind applied ethics. For instance, if we want to know whether it is okay to steal when hungry, we can look to different normative ethical theories as a way of providing an answer. So applied ethics rests to some extent on normative ethics.

Normative ethics is concerned with the development of ethical theories. Ethical theories aim to provide guidance as to *what* is right and wrong, as well as provide an explanation as to *why* it is so. What's more, normative ethical theories work at a basic level; that is, they are not simply a collection of principles or ideas, but a systematized account of what makes something right or wrong *fundamentally*.

Normative ethics will often draw on applied ethics, or at least on the kinds of concrete cases discussed by applied ethicists. For instance, many of the chapters in this book criticize ethical theories using thought experiments. In these experiments, a particular theory is applied to a concrete scenario. Then the critic draws attention to the fact that the theory's recommendation is striking or absurd in some way. However, fundamentally this book is not about applied ethics (or metaethics, the main exceptions being in the last section of the book, where we look at ethical views that reject generalization and theorizing in ethics). The other two levels of moral philosophy only arise here in the context of developing, defending, or criticizing systematic accounts of morality.

So why discuss normative ethical theory? Will studying ethics help you to be a better person? One way of thinking about the study of ethics is to analogize it to the study of fungi or astrophysics: you don't study fungi or astrophysics in order to be a better mushroom or star: mushrooms and stars are the *subject* of your inquiry. Similarly, according to this view, the job of ethicists is to understand what ethics *is*, rather than how to *be* more ethical ourselves.

On the other hand, ethics is clearly a subject closely tied to our practices. It discusses issues such as what we ought to value, how we ought to live, and how we should treat those around us. In this sense, ethics is both highly personal and directly applicable to our lives. Perhaps, then, ethics is more like subjects such as music or literature. We might reasonably expect that studying music and literature allows us to appreciate music or literature more fully, and perhaps even to *make* better music or literature; similarly, studying ethics can give us the conceptual resources for thinking about what you ought to value in your own life and what kind of person you should be: to appreciate our moral natures more fully.

Thinking about normative ethics also helps us to move past the obvious. Most people already have a sense of right and wrong. But in coming to appreciate the underlying nature of morality, we can understand not just *that* certain things are right or wrong, but *why* they are. And in coming to understand the *why* of ethics, we might also come to *change* our views about what things are right and wrong; we might find that our previous beliefs about morality were not well justified. Insofar as we care that our moral judgments are accurate—rather than merely expressing what we have unreflectively learned from others—engaging in ethical theorizing is an important and valuable activity.

Evaluating Moral Theories

We can distinguish between philosophical ethics—which is normative—and what we might call “sociological” or “anthropological” ethics—which is descriptive. If we were to take a sociological interest in ethics, our job would be merely to describe what beliefs people in different societies happen to hold regarding moral matters. We would note what they take to be valuable, what actions they take to be right or wrong, what kind of people they find morally admirable.

But philosophical ethics is not just about *what* people believe; it is more centrally about what people *ought to* believe. So, whereas a sociologist would describe a group’s beliefs, and then try to explain *how* those particular beliefs came about in the first place—perhaps in terms of social forces, environmental factors,

and so on—a philosopher is interested in what *arguments* there are in favor of those beliefs. The central concern of philosophers is therefore not the factors that *caused* certain beliefs, but whether those beliefs are *well justified*.

The purpose of this book is therefore not merely to present a range of ethical theories side by side and leave things there. Instead, the aim is to put these different theories into dialogue, with the goal of thinking through what the *right* theory of ethics might be. This textbook is therefore intended to facilitate a *conversation* between different proposals for the right ethical theory. The reader's task is to adjudicate this conversation; to assess and critically evaluate the different theories presented, with the goal of working out what the reader takes to be the most plausible of these theories, and why.

To establish whether any of the theories in this book are well justified, we need to adopt a critical stance toward them. In everyday terms, the word "criticism" typically has a negative valence. But in philosophy, criticism is a constructive, rather than destructive, activity. Productive criticism first involves trying to "think along" with the theories, to understand why someone might find them persuasive, what plausible justifications they might have. Then the reader can begin to evaluate the justifications offered for the theories and assess their plausibility. Ultimately, my hope is that the reader will be able to form or clarify their own position regarding normative ethics through a critical engagement with the ethical theories found in this book.

But how do we go about assessing the plausibility of moral theories? Moral philosophy is not necessarily difficult, but it has been going on for a long time, and in a short introduction like this one we do not have room to cover every possible factor that may come into play when evaluating a moral theory. However, below are a range of criteria that can get us started.

First, we can ask whether an ethical theory is internally consistent. If the various principles of the theory cannot be cashed out in logically coherent way, or the theory tells us that a particular action is both morally right *and* morally wrong at the same time, then we have very good reason to reject it, or at least to expect that it be modified to remove the contradiction.

We can also “test” moral theories. This is analogous to how we test scientific theories: we work out what the theory tells us should happen in novel cases—our prediction—then we run an experiment to see if what *actually* happens is consistent with our prediction. If so, this is evidence in favor of the theory; if not, it counts against it. While moral theories do not make predictions in quite the same way, we can do something similar: we can work out what the theory says we should do in various concrete cases and see how these “predictions” accord with our considered moral judgments. For instance, if a “thought experiment” tells us that a particular theory entails that genocide, mass murder, slavery, and rape are morally acceptable, this would count against the theory.

Of course, we must be careful in approaching ethics in this way. After all, our moral judgments are not foolproof, and it would be very surprising—and rather suspicious—if normative ethics did nothing more than confirm our society’s moral common sense. We therefore need to be open to the possibility that a moral theory might contradict what we take to be right or wrong, and that the mistake is with *us* rather than with the theory.

To apply this critical tool correctly, then, it helps to distinguish between levels of certainty and consensus and be aware of the ways in which our own beliefs may be parochial. For instance, our level of confidence about the wrongness of rape, genocide, or slavery is presumably much greater than our confidence about issues such as euthanasia, polyamory, or genetic engineering. If a moral theory disagrees with our beliefs about the latter, we should be far less sure that it is the theory that is wrong rather than us. We should also try to avoid drawing large conclusions from single cases. A moral theory that makes predictions at odds with our considered judgments across a *range* of cases is more likely to be wrong than one that gives us the “wrong” answer in just a few instances.

A good moral theory should also explain our wide-ranging practices of morality. Here this is not a matter of particular *judgments* of right and wrong, but rather the moral *concepts* we use. For instance, our everyday moral practices draw on ideas such as fairness, harm, loyalty, happiness, and so on. A good moral theory can make sense of the range of moral concepts we intuitively rely on, whereas a weak one has nothing to say about them or cannot

account for their importance. As with singular judgments, moral theorizing might mean rethinking or modifying some of our everyday moral concepts. But all things being equal, the more moral concepts a theory can incorporate or explain, the better.

Relatedly, the theory should be able to provide a plausible explanation for *why* certain things are moral or immoral. Does the justification the theory offers for why murder is wrong, for instance, make sense? Or does it explain the wrongness of murder by citing features that do not seem to really connect to what we take to be morally significant about it? For instance, one theory might explain the wrongness of murder in terms of the valuable future that has been taken from the victim, whereas another might understand the issue in terms of “theft,” that is, that the victim owns their own body, which the murderer then takes from them. Arguably, the “loss of a valuable future” explanation seems to capture the wrongness of murder more credibly than the “theft” account. If so, then the theory that gives rise to the “loss of a valuable future” view is the more plausible one.

It also counts in favor of a theory if it is compatible with a plausible moral psychology. That is, it can explain why we ought to be moral, why the fact that something is moral ought to motivate us.

Finally, a theory usually also gains support if it can *unify* our moral judgments. The more basic concepts a theory needs to rely on to account for our moral practices, the less plausible it seems. This is like the principle of parsimony in science: if more than one theory explains the data equally well, then we ought to prefer the simpler theory. Of course, there is always a balancing act to be had here with explanatory power, and the simpler theory may not always be better. But all other things being equal, we would usually say that a more parsimonious theory is more plausible than a less parsimonious one.

The Structure of the Book

This book differs from typical ethical theory textbooks in respect to the theories covered. Traditionally, the options presented are usually taken mostly—and often exclusively—from the Western philosoph-

ical canon. This book aims for a wider conversation and considers ethical theories from Western, East Asian, South Asian, and African philosophical traditions. When it comes to moral philosophy, this is no means exhaustive: to even attempt a truly comprehensive discussion of ethical theory from a truly “global” perspective would take up far more space than is available to us here. Given the huge diversity of fascinating and worthwhile ethical theories that could have been included, difficult choices had to be made.

On what basis, then, have I chosen the ethical theories in this book? My central aim with this book was to put the different theories we discuss into *conversation*. This means choosing theories that I am confident can speak to each other productively. But this leads to a familiarity bias, and the choice of theory therefore speaks at least partly to my own personal background and expertise. As such, a theory not being included in this book is in no way an indication that it is not worth taking seriously. So, while every theory we discuss in this book deserves its place in it, it’s important to note at the start the partial—and at least somewhat arbitrary—nature of the selection process.

Nonetheless, every chapter of this book provides an important, unique perspective on ethics from a range of the most significant, impactful, and long-lasting schools of thought in world culture. The reader should therefore hopefully gain an understanding of, and appreciation for, many of the key ethical theories that have shaped society throughout human history.

Structurally, I have divided the different theories found in the book into four categories: character-based theories, consequence-based theories, principle-based (or deontological) theories, and anti-theoretical/particularist/other theories.

The categories of the book are the fundamental *kind* of theory in Western analytic philosophy. Each kind of theory can contain a range of different views that disagree with each other in important ways, but all are united in terms of what is of *fundamental importance* in ethics. Character-based theories, for instance, take the development of a certain character to be what ethics is about, whereas consequentialist theories hold that at the most basic level we ought to be concerned with the consequences of our actions.

However, in our everyday moral decision-making, all the different approaches often seem to be saying the same thing. After

all, most commonsensically morally good actions reflect a virtuous character, lead to good consequences, *and* follow plausible moral principles all at the same time. For instance, robbing a bank is not the kind of thing a virtuous person would typically do; it leads to bad consequences (on the whole) and contravenes commonsense moral principles such as “do not steal” (as well as more fundamental ones such as “treat others with respect”).

Nonetheless, there are times—in principle at least, though almost certainly in practice as well—when character, consequences, and principles come apart. Sometimes acting from a virtuous character will lead to less-than-optimal consequences (why spend time and energy developing our moral character when children are starving, after all) or will require us to contravene an important moral principle (it may be kind to tell our friend that their new shirt looks good even when this requires us break a prohibition against lying). Sometimes achieving the best consequences requires acting in unvirtuous ways (for instance, killing an infant Hitler in his crib may lead to the best consequences, but it may not be the action of a truly virtuous person and certainly contravenes a prohibition against killing). And following the right moral principles may sometimes require us to forgo the optimal consequences (for instance, it may be better overall to harvest the organs of strangers against their will, but this certainly contravenes the principle of autonomy) or act in an unvirtuous way (coldly following a moral rule even when it lacks kindness).

Hence, in principle different approaches to ethics come apart at a basic level. For all the times that they converge in what they require, there are nonetheless times when they demand different—and fundamentally incompatible—actions. So, while it is tempting to subscribe to all the different approaches—to say “it’s a little bit of all of them”—this answer is just too easy. It certainly does seem that each approach offers something valuable. But if each approach is “grounded” in incompatible ways—if each approach sees a different thing as being that which *makes* something morally good or bad—then we need to be sensitive to these fundamental disagreements.

In saying this, any plausible moral theory almost certainly needs to incorporate character, consequences, and principles in *some* way. After all, these *all* seem central to our everyday moral

thinking. And most theories *do* incorporate all in some way. For instance, consequentialists value both principles and character. After all, even if we are concerned with bringing about the best consequences, it often is not easy to know what actions will in fact lead to the best consequences. Hence, adopting moral principles—such as “do not kill” or “do not steal”—as “rules of thumb” is likely to help us to act in beneficial ways. Similarly, developing certain traits of character—kindness, courage, patience—is in practice likely to lead to good consequences.

But note what is going on here: what *truly* matters in this picture are the consequences. Principles and character *only* matter insofar as they enable us to bring about desirable consequences. They do not have any value in and of themselves: their value is derivative rather than intrinsic. If a virtue or principle leads to *undesirable* consequences, for the consequentialist it ceases to have value.

Similarly, deontologists will probably need a plausible moral psychology. After all, being the kind of person who always follows the right principles is not easy. Developing our character in certain ways can help with this. So providing an account of how we connect up our basic human psychology to the principles we are obliged to follow is important. Similarly, consequences usually matter in at least *some* way to deontologists. For instance, respecting others will often involve intentionally acting in ways that benefit them. It is questionable whether we can describe someone as respecting others if they have no regard for how their actions impact those others. In fact, some deontological theories—such as contractualism and discourse ethics—build a concern for consequences into the generation of their moral principles in the first place.

But again, we need to be clear about what is being said here: for deontologists, the *fundamental* thing is abiding by the right principles. Character and consequences only play a role insofar as they help us to apply (or generate) those principles. The difference between the different theories is therefore in which feature *grounds* ethics and which operate at a higher, less fundamental level. Part of our task in ethical theorizing is therefore to be clear about the relationship between these features of ethics.

In short, *at the most basic level*, each of these approaches differs: each grounds morality in a different way. In this respect, they are

rivals, and we cannot agree with more than one. When it comes to higher-level matters—concrete questions of applied ethics—the different approaches may converge. Similarly, they may rest on (or be compatible with) the same metaethical views. But when it comes to the question of *what really matters* in ethics, these theories conflict with one another. Part of our task when reading this book is to be attentive to precisely how and why this is the case.

Global Philosophy

This book is divided into four sections, each containing a “kind” of ethical theory. Each of these sections contains ethical theories from different philosophical traditions. However, this system of categorization itself is one that developed within Western analytic philosophy and arguably reflects the interests and concerns of that tradition. It is questionable whether this way of dividing the theories truly captures the uniqueness of the different theories we discuss; in some ways it may distort them, forcing them to fit a shape for which they were not intended, or requiring them to speak to issues that matter to Western analytic philosophers but are not of central concern to the tradition they are from. Nonetheless, I have divided these theories in this way. Why?

There are two main reasons, both of which are about ensuring that this book achieves its aim of not merely *informing* the reader of different viewpoints, but rather enables a conversation between them: that it makes possible the activity of *global* philosophy.

The first is that conversations about grounding ethical theories in terms of character, consequences, and principles are already taking place in communities of philosophers who specialize in non-Western philosophical traditions. This division therefore may not be perfect, but it has been *useful enough* to provide a tool for thinking about ethics across a range of traditions. For instance, this book contains three different chapters on Buddhist ethics, each of which is found in a different section. While each chapter is nominally discussing a different branch of the extremely diverse Buddhist tradition, it also reflects an ongoing discussion among scholars of Buddhist philosophy about just what kind of moral theory Buddhism *really* advocates.

Similarly, there are two chapters on Confucianism (the first arguing that it is about character, the second that it is about principles). In fact, Confucian philosophers have discussed how Confucianism ought to be grounded at length, and there are advocates for placing it in each of the four sections of this book. The same kind of thing can be said about African philosophy: the chapter on Akan ethics outlines the extensive disagreements among Twi-speaking philosophers regarding how Akan ethics should be categorized according to the typology being used in this book.

In short, if the aim of this book is to enable conversation between ethical theories from different philosophical traditions, there are worse places to start than by using a system of categorization that—despite its shortcomings—has found widespread use among Western and non-Western philosophers alike.

Second, this categorization is a useful tool for us to begin to explore similarities and differences between different views. After all, without *some* common ground, conversation is impossible, and the various views discussed in this book will end up merely talking past one another. Whether the conversation will *remain* on the similarities and differences that this system of categorization helps us to identify—or whether the conversation will move on to different questions and concerns entirely—is very much an open question. But as a place to begin, dividing the theories in this way can be productive.

We therefore do not need to treat this system of classification as an immutable structure for the conversation to nevertheless find it useful as a tool to provoke discussion, debate, and disagreement by offering a starting point for engagement. And so long as the book's classificatory schema is taken in the provisional way that it is intended, it can be extremely useful.