# The Right Thing to Do

Readings in Moral Philosophy

EDITED BY
JAMES RACHELS
AND STUART RACHELS

EIGHTH
EDITION

The Right Thing to Do

The Right Thing to Do

Readings in Moral Philosophy

Eighth Edition

Edited by
James Rachels and Stuart Rachels

Contents

Preface

Moral philosophy is the study of how one should live. This anthology introduces that great subject. The readings cover the main moral theories and present a wealth of ideas about various practical matters.

In selecting the pieces for this volume, I looked for articles on serious moral topics that are deftly argued and pleasant to read, that lend themselves to lively discussion, and that reward careful study. I believe that the selections chosen are not merely good articles on suitable topics; they are first-rate essays on compelling issues. Students who read this book will want to read more, I think, unless the subject is simply not for them.

This eighth edition contains eight new essays, on eight new topics: the opioid crisis, hazing, the meaning of life, the ethics of taxation, guns, doping in sports, end-of-life care, and the history of racism in America. Seven selections from the seventh edition were cut: the essays on organ-selling and eugenics (because instructors made less use of them), the pieces by Martin Luther King Jr. and Douglas Lackey (due to high permissions fees), and the selections on same-sex marriage (because that issue has been satisfactorily resolved in American law, in a way that most college-age Americans support).

For their help in preparing the new edition, I thank David Connelly, Glenda Elliott, Heather Elliott, Sean McAleer, Jake Osachy, Jada Posey, and Carol Rachels. I also thank my new publisher—Rowman & Littlefield—and in particular Natalie Mandziuk.

To learn more about James Rachels, visit www.jamesrachels.org.

If you have suggestions for the next edition, please let me know!

—Stuart Rachels

INTRODUCTION

A Short Introduction to Moral Philosophy

James Rachels

An ancient legend tells the story of Gyges, a poor shepherd who found a magic ring in a fissure opened by an earthquake. The ring would make its wearer invisible, so he could go anywhere and do anything undetected. Gyges was an unscrupulous fellow, and he quickly realized that the ring could be put to good advantage. We are told that he used its power to gain entry to the royal palace, where he seduced the queen, murdered the king, and seized the throne. (It is not explained how invisibility helped him to seduce the queen—but let that pass.) In no time at all, he went from being a poor shepherd to being king of all the land.

This story is recounted in Book II of Plato's *Republic*. Like all of Plato's works, *The Republic* is written in the form of a dialogue between Socrates and his companions. Glaucon, who is having an argument with Socrates, uses the story of Gyges's ring to make a point.

Glaucon asks us to imagine that there are two such rings, one given to a man of virtue and the other given to a rogue. How might we expect them to behave? The rogue, of course, will do anything necessary to increase his own wealth and power. Since the cloak of invisibility will protect him from discovery, he can do anything he pleases without fear of being caught. Therefore, he will recognize no moral constraints on his conduct, and there will be no end to the mischief he will do.

But how will the so-called virtuous man behave? Glaucon suggests that he will do no better than the rogue:

> No one, it is commonly believed, would have such iron strength of mind as to stand fast in doing right or keep his hands off other men's goods, when he could go to the marketplace and fearlessly help himself to anything he wanted, enter houses and sleep with any woman he chose, set prisoners free and kill men at his pleasure, and in a word go about among men with the powers of a god. He would behave no better than the other; both would take the same course.

Moreover, Glaucon asks, why shouldn't he? Once he is freed from the fear of reprisal, why shouldn't a person simply do what he pleases, or what he thinks is best for himself? Why should he care at all about "morality"?

The Republic, written more than 2,300 years ago, was one of the first great works of moral philosophy in Western history. Since then, philosophers have formulated theories to explain what morality is, why it is important, and why it has the peculiar hold on us that it does. What, if anything, justifies our belief that we *morally ought* to act in one way rather than another?

RELATIVISM

Perhaps the oldest philosophical theory about morality is that right and wrong are relative to the customs of one's society—on this view, there is nothing behind the demands of morality except social convention. Herodotus, the first of the great Greek historians, lived at about the time of Socrates. His *History* is full of wonderful anecdotes that illustrate his belief that "right" and "wrong" are little more than names for social conventions. Of the Massagetae, a tribe in Central Asia, he writes:

> The following are some of their customs—Each man has but one wife, yet all the wives are held in common. . . . Human life does not come to its natural close with these people; but when a man grows very old, all his kinsfolk collect together and offer him up in sacrifice; offering at the same time some cattle also. After the sacrifice they boil the flesh and feast on it; and those who thus end their days are reckoned the happiest. If a man dies of disease they do not eat him, but bury him in the ground, bewailing his ill-fortune that he did not come to be sacrificed. They sow no grain, but live on their herds, and on fish, of which there is great plenty in the Araxes. Milk is what they chiefly drink. The only god they worship is the

sun, and to him they offer the horse in sacrifice, under the notion of giving the swiftest of the gods the swiftest of all mortal creatures.

Herodotus did not think the Massagetae were to be criticized for such practices. Their customs were neither better nor worse than those of other peoples; they were merely different. The Greeks, who considered themselves more "civilized," might have thought that their customs were superior, but, Herodotus says, that is only because everyone believes the customs of his own society to be the best. The "truth" depends on one's point of view—that is, on the society in which one happens to have been raised.

Relativists think that Herodotus was obviously on to something and that those who believe in "objective" right and wrong are merely naive. Critics, however, object to the theory on a number of grounds. First, it is exceedingly conservative, in that the theory endorses whatever moral views happen to be current in a society. Consider our own society. Many people believe that our society's moral code is mistaken, at least on some points—for example, they may disagree with the dominant social view regarding capital punishment or gay rights or the treatment of nonhuman animals. Must we conclude that these would-be reformers are wrong, merely because they oppose the majority view? Why must the majority always be right?

But there is a deeper problem with Relativism, emphasized by Socrates. Some social customs are, indeed, merely arbitrary, and when these customs are at issue it is fruitless to insist that one of society's practices are better than another's. Funerary practices are a good example. The Greeks burned their dead, while the Callatians ate their dead, but neither practice is better than the other. However, it does not follow from this that *all* social practices are arbitrary in the same way. Some are, and some are not. The Greeks and the Callatians were free to accept whatever funerary practices they liked because no objective reason could be given why one practice was superior to the other. In the case of other practices, however, there may be good reasons why some are superior. It is not hard, for example, to explain why honesty and respect for human life are socially desirable, and similarly it is not hard to explain why slavery and racism are undesirable. Because we can support our judgments about these matters with rational arguments, we do not have to regard those judgments as "merely" the expression of our particular society's moral code.

DIVINE COMMANDS

A second ancient idea, also familiar to Socrates, was that moral living consists in obedience to divine commands. If this were true, then we could easily answer the challenge of Gyges's ring—even if we had the power of invisibility, we would still be subject to divine retribution, so ultimately we could not "get away with" doing whatever we wanted.

But Socrates did not believe that right living could consist merely in trying to please the gods. In the *Euthyphro*, another of Plato's dialogues, Socrates is shown considering at some length whether "right" can be the same as "what the gods command." Now we may notice, to begin with, that there are considerable practical difficulties with this as a general theory of ethics. How, for example, are we supposed to *know* what the gods command? There are those who claim to have spoken with God about the matter and who therefore claim to be in a position to pass on his instructions to the rest of us. But people who claim to speak for God are not the most trustworthy folks—hearing voices can be a sign of schizophrenia or megalomania just as easily as an instance of divine communication. Others, more modestly, rely on scripture or church tradition for guidance. But those sources are notoriously ambiguous—they give vague and often contradictory instructions—so, when people consult these authorities, they typically rely on whatever elements of scripture or church tradition support the moral views they are already inclined to agree with. Moreover, because scripture and church tradition have been handed down from earlier times, they provide little direct help in addressing distinctively contemporary problems: the problem of environmental preservation, for example, or the problem of how much of our resources should be allocated to cancer research as opposed to other worthy endeavors.

Still, it may be thought that God's commands provide the ultimate *authority* for ethics, and that is the issue Socrates addressed. Socrates accepted that the gods exist and that they may issue instructions. But he showed that this cannot be the ultimate basis of ethics. He pointed out that we have to distinguish two possibilities: Either the gods have some reason for the instructions they issue, or they do not. If they do not, then their commands are merely arbitrary—the gods are like petty tyrants who demand that we act in this way and that, even though there is no good reason for it. But this is an impious view that religious people will

not want to accept. On the other hand, if we say that the gods do have good reasons for their instructions, then we have admitted that there is a standard of rightness independent of their commands—namely, the standard to which the gods themselves refer in deciding what to require of us.

It follows, then, that even if one accepts a religious picture of the world, the rightness or wrongness of actions cannot be understood merely in terms of their conformity to divine prescriptions. We may always ask why the gods command what they do, and the answer to *that* question will reveal why right actions are right and why wrong actions are wrong.

ARISTOTLE

Although Relativism and the Divine Command Theory have always had supporters, they have never been popular among serious students of moral philosophy. The first extended, systematic treatise on moral philosophy, produced two generations after Socrates, was Aristotle's *Nicomachean Ethics* (ca. 330 BC), and Aristotle wasted no time on such notions. Instead, Aristotle offered a detailed account of the virtues—the qualities of character that people need to do well in life. The virtues include courage, prudence, generosity, honesty, and many more; Aristotle sought to explain what each one is and why it is important. His answer to the question of Gyges's ring was that virtue is necessary for human beings to achieve happiness; therefore, the man of virtue is ultimately better off *because* he is virtuous.

Aristotle's view of the virtuous life was connected with his overall way of understanding the world and our place in it. Aristotle's conception of what the world is like was enormously influential; it dominated Western thinking for more than 1,700 years. A central feature of this conception was that *everything in nature exists for a purpose.* "Nature," Aristotle said, "belongs to the class of causes which act for the sake of something."

It seems obvious that artifacts such as knives and chariots have purposes, because we have their purposes in mind when we make them. But what about natural objects that we do not make? Do they have purposes too? Aristotle thought so. One of his examples was that we

have teeth so that we can chew. Such biological examples are quite per-
suasive; the parts of our bodies do seem, intuitively, to have particular
purposes—eyes are for seeing, the heart is for pumping blood, and so
on. But Aristotle's thesis was not limited to organic beings. According
to him, *everything* in nature has a purpose. He also thought, to take a
different sort of example, that rain falls so that plants can grow. As
odd as it may seem to a modern reader, Aristotle was perfectly serious
about this. He considered other alternatives, such as that the rain falls
"of necessity" and that this helps the plants only "by coincidence," and
rejected them. His considered view was that plants and animals are
what they are, and that the rain falls as it does, "because it is better so."

The world, therefore, is an orderly, rational system, with each thing
having its own proper place and serving its own special purpose. There
is a neat hierarchy: The rain exists for the sake of the plants, the plants
exist for the sake of the animals, and the animals exist—of course—for
the sake of people. Aristotle says: "If then we are right in believing
that nature makes nothing without some end in view, nothing to no
purpose, it must be that nature has made all things specifically for the
sake of man." This worldview is stunningly anthropocentric, or human-
centered. But Aristotle was hardly alone in having such thoughts;
almost every important thinker in human history has advanced such a
thesis. Humans are a remarkably vain species.

NATURAL LAW

The Christian thinkers who came later found Aristotle's view of the
world appealing. There was only one thing missing: God. Thus, the
Christian thinkers said that the rain falls to help the plants because
that is what the Creator intended, and the animals are for human use
because *that is what God made them for*. Values and purposes were,
therefore, conceived to be a fundamental part of the nature of things,
because the world was believed to have been created according to a
divine plan.

This view of the world had a number of consequences for ethics. On
the most general level, it affirmed the supreme value of human life,
and it explained why humans are entitled to do whatever they please
with the rest of nature. The basic moral arrangement—human beings,

whose lives are sacred, dominating a world made for their benefit—was enshrined as the Natural Order of Things.

At a more detailed level, a corollary of this outlook was that the "laws of nature" specify how things *ought to be*, as well as describing how things *are*. In turn, knowing how things ought to be enables us to evaluate states of affairs as objectively good or bad. Things are as they ought to be when they are serving their natural purposes; when they do not or cannot serve those purposes, things have gone wrong. Thus, teeth that have decayed and cannot be used for chewing are defective; and drought, which deprives plants of the rain they need, is a natural, objective evil.

There are also implications for human action: On this view, moral rules are one type of law of nature. The key idea here is that some forms of human behavior are "natural" while others are not; and "unnatural" acts are said to be wrong. Beneficence, for example, is natural for us because God has made us as social creatures. We want and need the friendship of other people, and we have natural affections for them; hence, behaving brutishly toward them is unnatural. Or to take a different sort of example, the purpose of the sex organs is procreation. Thus, any use of them for other purposes is "contrary to nature"—which is why the Christian church has traditionally regarded any form of sexual activity that cannot result in pregnancy, such as masturbation, gay sex, or sex with contraceptives, as impermissible.

This combination of ideas, together with others like them, formed the core of an outlook known as natural-law ethics. The Theory of Natural Law was developed most fully by Saint Thomas Aquinas (1225–1274), who lived at a time when the Aristotelian worldview was unchallenged. Aquinas was the foremost thinker among traditional Catholic theologians. Today natural-law theory still has adherents inside the Catholic Church, but few outside. The reason is that the Aristotelian worldview, on which natural-law ethics depends, has been replaced by the outlook of modern science.

Galileo, Newton, Darwin, and others developed ways of understanding natural phenomena that made no use of evaluative notions. In their way of thinking, the rain has no purpose. It does not fall in order to help the plants grow. Plants typically get the amount of water they need because each species has evolved, by natural selection, in the environment in which that amount of water is available. Natural selection

produces an orderly arrangement that *appears* to have been designed, but that is only an illusion. To explain nature there is no need to assume purpose-involving principles, as Aristotle and the Christians had done. This new outlook was threatening to the Catholic Church, and they condemned it.

Modern science transformed people's view of what the world is like. But part of the transformation, inseparable from the rest, was an altered view of the nature of ethics. Right and wrong could no longer be deduced from the nature of things, for on the new view the natural world does not, in and of itself, manifest value and purpose. The *inhabitants* of the world may have needs and desires that generate values special to them, but that is all. The world apart from those inhabitants knows and cares nothing for their values, and it has no values of its own. A hundred and fifty years before Nietzsche declared, "There are no moral facts," the Scottish philosopher David Hume had come to the same conclusion. Hume summed up the moral implications of the new worldview in his *Treatise of Human Nature* (1739) when he wrote:

> Take any action allow'd to be vicious: Willful murder, for instance. Examine it in all lights, and see if you can find that matter of fact, or real existence, which you call *vice.* In whichever way you take it, you find only certain passions, motives, volitions and thoughts. There is no other matter of fact in the case.

To Aristotle's idea that "nature has made all things for the sake of man," Hume replied: "The life of a man is of no greater importance to the universe than that of an oyster."

THE SOCIAL CONTRACT

If morality cannot be based on God's commands, nor on the idea of natural purpose, then what can it be based on? Ethics must somehow be understood as a purely human phenomenon—as the product of human needs, interests, and desires—and nothing else. Figuring out how to do this has been the basic project of moral philosophy from the seventeenth century on.

Thomas Hobbes, the foremost English philosopher of the seventeenth century, suggested one way in which ethics might be understood in

purely human terms. Hobbes assumed that "good" and "bad" are just names we give to things we like and dislike. Thus, because we may like different things, we may disagree about what is good or bad. However, Hobbes said, in our fundamental psychological makeup we are all very much alike. We are all basically self-interested creatures who want to live and to live as well as possible. This is the key to understanding ethics. Ethics arises when people realize *what they must do* to live well.

Hobbes was the first important modern thinker to provide a secular, naturalistic basis for ethics. He pointed out that each of us is enormously better off living in a mutually cooperative society than we would be if we tried to make it on our own. The benefits of social living go far beyond companionship: Social cooperation makes it possible to have schools, hospitals, and highways; houses with electricity and central heating; airplanes and telephones; websites and books; movies, opera, and bingo; science and agriculture. Without social cooperation we would lose these benefits and more. Therefore, it is to the advantage of each of us to do whatever is necessary to establish and maintain a cooperative society.

But it turns out that a mutually cooperative society can exist only if we adopt certain rules of behavior—rules that require telling the truth, keeping our promises, respecting one another's lives and property, and so on:

Without the presumption that people will tell the truth, there would be no reason for people to pay any attention to what other people say. Communication would be impossible. And without communication among its members, society would collapse.

Without the requirement that people keep their promises, there could be no division of labor—workers could not count on getting paid, retailers could not rely on their agreements with suppliers, and so on—and the economy would collapse. There could be no business, no building, no agriculture, no medicine.

Without assurances against assault, murder, and theft, no one could feel secure; everyone would have to be constantly on guard against everyone else, and social cooperation would be impossible.

Thus, to obtain the benefits of social living, we must strike a bargain with one another, with each of us agreeing to obey these rules, provided others do likewise. We must also establish mechanisms for enforcing

these rules—such as legal sanctions and other, less formal methods of enforcement—so that we can *count on* one another to obey them. This "social contract" is the basis of morality. Indeed, morality can be defined as nothing more or less than *the set of rules that rational people will agree to obey, for their mutual benefit, provided that other people will obey them as well*.

This way of understanding morality has a number of appealing features. First, it takes the mystery out of ethics and makes it a practical, down-to-earth business. Living morally is not a matter of blind obedience to the mysterious dictates of a supernatural being; nor is it a matter of fidelity to lofty but pointless abstract rules. Instead, it is a matter of doing what it takes to make social living possible.

Second, the Social Contract Theory explains why we should *care* about ethics—it offers at least a partial response to the problem of Gyges's ring. If there is no God to punish us, why should we bother to do what is "right," especially when it is not to our advantage? The answer is that it *is* to our advantage to live in a society where people behave morally—thus, it is rational for us to accept moral restrictions on our conduct as part of a bargain we make with other people. We benefit directly from the ethical conduct of others, and our own compliance with the moral rules is the price we pay to secure their compliance.

Third, the Social Contract approach gives us a sensible and mature way of determining what our ethical duties really are. When "morality" is mentioned, the first thing that pops into many people's minds is an attempt to restrict their sex lives. It is unfortunate that the word *morals* has come to have such a connotation. The whole purpose of having a system of morality, according to Social Contract Theory, is to make it possible for people to live their individual lives in a setting of social cooperation. Its purpose is *not* to tell people what kinds of lives they should lead (except insofar as it is necessary to restrict conduct in the interests of maintaining social cooperation). Therefore, an ethic based on the Social Contract would have little interest in what people do in their bedrooms.

Finally, we may note again that the Social Contract Theory assumes relatively little about human nature. It treats human beings as self-interested creatures and does not assume that they are naturally altruistic, even to the slightest degree. One of the theory's charms is that it can reach the conclusion that we ought, often, to *behave* altruistically,

without assuming that we *are* naturally altruistic. We want to live as well as possible, and moral obligations are created as we band together with other people to form the cooperative societies that are necessary for us to achieve this fundamentally self-interested goal.

ALTRUISM AND SELF-INTEREST

Are people essentially self-interested? Although the Social Contract Theory continues to attract supporters, not many philosophers and psychologists today would accept Hobbes's egoistic view of human nature. It seems evident that humans have at least *some* altruistic feelings, if only for their family and friends. We have evolved as social creatures just as surely as we have evolved as creatures with legs—caring for our kin and members of our local group is as natural for us as walking.

If humans do have some degree of natural altruism, does this have any significance for morals? Hume thought so. Hume agreed with Hobbes that our moral opinions are expressions of our feelings. In 1739, when he invited his readers to consider "willful murder" and see if they could find that "matter of fact" called "vice," Hume concluded:

> You can never find it, till you turn your reflexion into your own breast, and find a sentiment of [disapproval], which arises in you, towards this action. Here is a matter of fact; but 'tis the object of feeling. . . . It lies in yourself, not in the object. So that when you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it.

And what, exactly, is "the constitution of our nature"? Of course, it is part of our nature to care about ourselves and our own welfare. But Hume added that we also have "*social* sentiments"—feelings that connect us with other people and make us concerned about their welfare. That is why, Hume said, we measure right and wrong by "the true interests of mankind":

> In all determinations of morality, this circumstance of public utility is ever principally in view; and wherever disputes arise, either in philosophy or common life, concerning the bounds of duty, the question cannot, by

any means, be decided with greater certainty than by ascertaining, on any side, the true interests of mankind.

This view came to be known as Utilitarianism. In modern moral philosophy, it is the chief alternative to the Social Contract Theory.

UTILITARIANISM

Utilitarians believe that one principle sums up all of our moral duties. The precept is that *we should always do whatever will produce the greatest possible balance of happiness over unhappiness for everyone who will be affected by our action*. This "principle of utility" is deceptively simple. It is actually a combination of three ideas: First, in determining what to do, we should be guided by the expected consequences of our actions—we should do whatever will have the best consequences. Second, in determining which consequences are best, we should give the greatest possible weight to the happiness or unhappiness that would be caused—we should do whatever will cause the most happiness or the least unhappiness. And finally, the principle of utility assumes that each individual's happiness is equally as important as anyone else's.

Although Hume expressed the basic idea of Utilitarianism, two other philosophers elaborated it in greater detail. Jeremy Bentham, an Englishman who lived in the late eighteenth and early nineteenth centuries, was the leader of a group of philosophical radicals who aimed to reform the laws of Britain along utilitarian lines. This group was remarkably successful in advancing such causes as prison reform and restrictions on the use of child labor. John Stuart Mill, the son of one of Bentham's original followers, gave the theory its most popular and influential defense in his book *Utilitarianism*, published in 1861.

The Utilitarian movement attracted critics from the start. It was an easy target because it ignored conventional religious notions. The point of morality, according to the Utilitarians, had nothing to do with obeying God or gaining credit in Heaven. Rather, the point was just to make life in this world as comfortable and happy as possible. So some critics condemned Utilitarianism as a godless doctrine. To this Mill replied:

[T]he question depends upon what idea we have formed of the moral character of the Deity. If it be a true belief that God desires, above all

things, the happiness of his creatures, and that this was his purpose in their creation, utility is not only not a godless doctrine, but more profoundly religious than any other.

Utilitarianism was also an easy target because it was (and still is) a *subversive* theory, in that it turned many traditional moral ideas upside down. Bentham argued, for example, that the purpose of the criminal justice system cannot be understood in the traditional way as "paying back" wrongdoers for their wicked deeds—that only piles misery upon misery. Instead, the social response to crime should be threefold: to identify and deal with the causes of criminal behavior; where possible, to reform individual lawbreakers and make them into productive citizens; and to "punish" people only insofar as it is necessary to deter others from committing similar crimes. Or, to take a different example, by insisting that everyone's happiness is equally important, the Utilitarians offended various elitist notions of group superiority. According to the Utilitarian standard, neither race, sex, nor social class makes a difference to one's moral status. Mill himself wrote a book on *The Subjection of Women* that became a classic of the nineteenth-century suffragist movement.

Finally, Utilitarianism was controversial because it had no use for "absolute" moral rules. The Utilitarians regarded the traditional rules—against killing, lying, breaking one's promises, and so on—as "rules of thumb," useful because following them will generally be for the best. But they are not absolute—whenever breaking such a rule will have better results for everyone concerned, the rule should be broken. The rule against killing, for example, might be suspended in the case of voluntary euthanasia for someone dying of a painful illness. Moreover, the Utilitarians regarded some traditional rules as dubious, even as rules of thumb. For example, Christian moralists had traditionally said that masturbation is evil because it violates the Natural Law; but from the point of view of the Principle of Utility, it appears to be harmless. A more serious matter is the traditional religious condemnation of homosexuality, which has resulted in misery for countless people. Utilitarianism implies that if an activity makes people happy, without harming anyone, it cannot be wrong.

But it is one thing to describe a moral view; it is another thing to justify it. Utilitarianism says that our moral duty is to "promote the general

happiness." Why should we do that? How can the challenge of Gyges's ring be answered? As Mill puts it:

> I feel that I am bound not to rob or murder, betray or deceive; but why am I bound to promote the general happiness? If my own happiness lies in something else, why may I not give that the preference?

Aside from the "external sanctions" of law and public opinion, Mill thinks there is only one possible reason for accepting this or any other moral standard. The "internal sanction" of morality must always be "a feeling in our minds," regardless of what sort of ethic this feeling endorses. The kind of morality we accept will, therefore, depend on the nature of our feelings: If human beings have "social feelings," then Mill says that utilitarian morality will be the natural standard for them:

> The firm foundation [of utilitarian morality] is that of the social feelings of mankind—the desire to be in unity with our fellow creatures, which is already a powerful principle in human nature, and happily one of those which tend to become stronger, even without express inculcation, from the influences of advancing civilization.

IMPARTIALITY

Utilitarianism, as we have seen, has implications that are at odds with traditional morality. Much the same could be said about the Social Contract Theory. In most of the practical matters that have been mentioned—punishment, racial discrimination, women's rights, euthanasia, gay rights—the two theories have similar implications. But there is one matter on which they differ dramatically. Utilitarians believe that we have a very extensive moral duty to help other people. Social Contract theorists deny this.

Suppose, for example, you are thinking of spending $1,000 for a new living room carpet. Should you do this? What are the alternatives? One alternative is to give the money to an agency such as the United Nations Children's Fund (UNICEF). Each year millions of third-world children die of easily preventable diseases because there isn't enough money to provide the vitamin-A capsules, antibiotics, and oral rehydration treatments they need. By giving the money to UNICEF, and making do

a while longer with your old carpet, you could provide much-needed medical care for dozens of children. From the point of view of utility—seeking the best overall outcome for everyone concerned—there is no doubt you should give the money to UNICEF. Obviously, the medicine will help the kids a lot more than the new rug will help you.

But from the point of view of the Social Contract, things look very different. If morality rests on an agreement between people—remember, an agreement they enter into *to promote their own interests*—what would the agreement say about helping other people? Certainly, we would want the contract to impose a duty not to harm other people, even strangers. Each of us would obviously benefit from that. And it might be in our best interests to accept a mutual obligation to provide aid to others when it is easy and convenient to do so. But would rational people accept a general duty to provide virtually unlimited aid to strangers, even at great cost to themselves? From the standpoint of self-interest, that sounds crazy. Jan Narveson, a contract theorist, writes in his book *Moral Matters* (1993):

> [M]orals, if they are to be rational, must amount to agreements among people—people of all kinds, each pursuing his or her own interests, which are various and do not necessarily include much concern for others and their interests. But people . . . have a broad repertoire of powers including some that can make them exceedingly dangerous, as well as others that can make them very helpful. This gives us reason to agree with each other that we will refrain from harming others in the pursuit of our interests, to respect each other's property and grant extensive civil rights, but not necessarily to go very far out of our way to be very helpful to those we don't know and may not particularly care for.

Unlike many philosophers who prefer to keep things abstract, Narveson is good about spelling out the implications of his view in a way that leaves no room for misunderstanding:

> What about parting with the means for making your sweet little daughter's birthday party a memorable one, in order to keep a dozen strangers alive on the other side of the world? Is this something you are morally required to do? Indeed not. She may well *matter* to you more than they. This illustrates again the fact that people do *not* "count equally" for most of us. Normal people care more about some people than others, and build their very lives around those carings.

Which view is correct? Do we have a moral duty to provide extensive aid to strangers, or not? Both views appeal ultimately to our emotions. A striking feature of Narveson's argument is its appeal to the fact that we *care more* for some people than others. This is certainly true: As he says, we care more for our own children than for "strangers on the other side of the world." But does this really mean that I may choose some trivial benefit for my children over the very lives of the strangers? Suppose there are two buttons on my desk at this moment, and by pressing button A, I can provide my son with a nice party; by pressing B, I can save the lives of a dozen strangers. Is it really all right for me to press A, just because I "care more" for my son? Mill agrees that the issue must be decided on the basis of feelings (how else could it be?), but for him it is not these small-scale personal feelings that have the final say. Instead, it is one's "conscientious feelings"—the feelings that prevail after everything has been thought through—that finally determine one's obligations. Mill assumes that we cannot, when we are thoughtful and reflective, approve of pushing button A.

However, some contemporary Utilitarians have argued that the matter need not be left to the uncertainties of individual feeling. It may be true, they say, that we all care more for ourselves, our family, and our friends than we care for strangers. But we have rational capacities as well as feelings, and if we think objectively about the matter, we will realize that other people are no different. Others, even strangers, also care about themselves, their families, and their friends, in the same way that we do. Their needs and interests are comparable to our own. In fact, *there is nothing of this general sort that makes anyone different from anyone else*—and if we are in all relevant respects similar to one another, then there is no justification for anyone taking his or her own interests to be more important. Peter Singer, a utilitarian philosopher, writes in his book *How Are We to Live?* (1995):

> Reason makes it possible for us to see ourselves in this way. . . . I am able to see that I am just one being among others, with interests and desires like others. I have a personal perspective on the world, from which my interests are at the front and center of the stage, the interests of my family and friends are close behind, and the interests of strangers are pushed to the back and sides. But reason enables me to see that others have similarly subjective perspectives, and that from "the point of view of the

universe" my perspective is no more privileged than theirs. Thus my ability to reason shows me the possibility of detaching myself from my own perspective, and shows me what the universe might look like if I had no personal perspective.

So, from an objective viewpoint, each of us must acknowledge that our own perspective—our own particular set of needs, interests, likes, and dislikes—is only one among many and has no special status.

KANT

The idea of impartiality is also central to the third major alternative in modern moral philosophy, the system of ethical ideas devised by the great German philosopher Immanuel Kant (1724–1804). Like the Social Contract theorists and the Utilitarians, Kant sought to explain ethics without appealing to divine commands or "moral facts." Kant's solution was to see morality as a product of "pure reason." Just as we must do some things because of our *desires*—for example, because I desire to go to a concert, I must buy a ticket—the moral law is binding on us because of our *reason*.

Like the Utilitarians, Kant believed that morality can be summed up in one ultimate principle, from which all our duties and obligations are derived. But his version of the "ultimate moral principle" was very different from the Principle of Utility, because Kant did not emphasize the outcomes of actions. What was important for him was "doing one's duty," and he held that a person's duty is not determined by calculating consequences.

Kant called his ultimate moral principle the "Categorical Imperative." But he gave this principle two very different formulations. The first version of the Categorical Imperative, as expressed in his *Fundamental Principles of the Metaphysics of Morals* (1785), goes like this:

> Act only according to that maxim by which you can at the same time will that it should become a universal law.

Stated in this way, Kant's principle summarizes a procedure for deciding whether an act is morally permissible. When you are contemplating

a particular action, you are to ask what rule you would be following if you were to do it. (This will be the "maxim" of the act.) Then you are to ask whether you would be willing for that rule to be followed by everyone all the time. (That would make it a "universal law" in the relevant sense.) If so, the rule may be followed, and the act is permissible. However, if you would not be willing for everyone to follow the rule, then you may not follow it, and the act is morally impermissible.

This explains why the Moral Law is binding on us simply by virtue of our rationality. The first requirement of rationality is that we be consistent, and it would not be consistent to act on a maxim that we could not want others to adopt as well. Kant believed, in addition, that consistency requires us to interpret moral rules as having no exceptions. For this reason, he endorsed a whole range of absolute prohibitions, covering everything from lying to suicide.

However, Kant also gave another formulation of the Categorical Imperative. Later in the same book, he said that the ultimate moral principle may be understood as saying:

> So act that you treat humanity, whether in your own person or in that of another, always as an end and never as means only.

What does it mean to say that persons are to be treated as "ends" and never as a "means only"? Kant gives this example: Suppose you need money, and so you want a "loan," but you know you could not repay it. In desperation, you consider making a false promise (to repay) in order to trick a friend into giving you the money. May you do this? Perhaps you need the money for a good purpose—so good, in fact, that you might convince yourself that the lie would be justified. Nevertheless, if you lied to your friend, you would merely be manipulating him and using him "as a means."

On the other hand, what would it be like to treat your friend "as an end"? Suppose you told the truth—that you need the money for a certain purpose but could not repay it. Then your friend could make up his own mind about whether to let you have it. He could exercise his own powers of reason, consulting his own values and wishes, and make a free, autonomous choice. If he did decide to give the money for this purpose, he would be choosing to make that purpose his own. Thus, you would not merely be using him as a means to achieving your goal.

CONCLUSION

Our purpose here is not to reach any firm conclusion about which of these approaches, if any, is correct. But we may end with an observation about how that project might be undertaken.

Philosophical ideas are often very abstract, and it is difficult to see what sort of evidence counts for or against them. It is easy enough to appreciate, intuitively, the ideas behind each of these theories, but how do we determine which, if any, is correct? It is a daunting question. Faced with this problem, people are tempted to accept or reject philosophical ideas on the basis of their intuitive appeal—if an idea sounds good, one may embrace it; or if it rubs one the wrong way, it may be discarded. But this is hardly a satisfactory way to proceed if we want to discover the truth. How an idea strikes us is not a reliable guide, for our "intuitions" may be mistaken.

Happily, there is an alternative. An idea is no better than the arguments that support it. So, to evaluate a philosophical idea, we may examine the reasoning behind it. The great philosophers knew this: They did not simply announce their philosophical opinions; instead, they presented arguments in support of their views. The leading idea, from the time of Socrates to the present, has been that truth is discovered by considering the reasons for and against the various alternatives. So the correct theory is the one supported by the best arguments. Thus, philosophical thinking consists, to a large extent, of formulating and assessing arguments. This is not the whole of philosophy, but it is a big part of it. It is what makes philosophy a rational enterprise.

Chapter 2

Some Basic Points about Arguments

James Rachels

Philosophy without argument would be a lifeless exercise. What good would it be to produce a theory if there were no reasons for thinking it correct? And of what interest is the rejection of a theory if there are no good reasons for thinking it incorrect? A philosophical idea is exactly as good as the arguments in its support.

Therefore, if we want to think clearly about philosophical matters, we have to learn something about the evaluation of arguments. We have to learn to distinguish the sound ones from the unsound ones. This can be a tedious business, but it is indispensable if we want to discover the truth.

ARGUMENTS

In ordinary English, the word *argument* often means a fight, and there is a hint of unpleasantness in the word. That is not the way the word is used here. In the logician's sense, an argument is a chain of reasoning designed to prove something. It consists of one or more *premises* and a *conclusion,* together with the claim that the conclusion *follows from* the premises. Here is a simple argument. This example is not particularly interesting in itself, but it is short and clear, and it will help us grasp the main points.

(1) All men are mortal.
Socrates is a man.
Therefore, Socrates is mortal.

The first two statements are the premises; the third statement is the conclusion; and the word *therefore* indicates that the conclusion is supposed to follow from the premises.

What does it mean to say that the conclusion "follows from" the premises? It means that a certain logical relation exists between the premises and the conclusion, namely, that *if* the premises are true, then the conclusion must be true also. (Another way to put the same point is: The conclusion follows from the premises if, and only if, it is impossible for the premises to be true and the conclusion false at the same time.) In example (1), we can see that the conclusion does follow from the premises. If it is true that all men are mortal, and Socrates is a man, then it must be true that Socrates is mortal. (Or, it is impossible for it to be true that all men are mortal, and for Socrates to be a man, and be false that Socrates is mortal.)

In example (1), the conclusion follows from the premises, *and* the premises are in fact true. However, the conclusion of an argument may follow from the premises even if the premises are not actually true. Consider this argument:

(2) All people from Georgia are famous.
    Jimmy Carter is from Georgia.
    Therefore, Jimmy Carter is famous.

Clearly, the conclusion of this argument does follow from the premises: *If* it were true that all Georgians are famous, and Jimmy Carter is from Georgia, then it would follow that Jimmy Carter is famous. This logical relation holds between the premises and conclusion even though one of the premises is in fact false.

At this point, logicians customarily introduce a bit of terminology. They say that an argument is *valid* just in case its conclusion follows from its premises. Both of the examples given above are valid arguments in this technical sense.

In order to be a *sound* argument, however, two things are necessary: The argument must be valid, *and* its premises must be true. Thus, the argument about Socrates is a sound argument, but the argument about Jimmy Carter is not sound, because even though it is valid, its premises are not all true.

It is important to notice that an argument may be unsound even though its premises and conclusion are both true. Consider the following silly example:

(3) The earth has one moon.
   John F. Kennedy was assassinated.
   Therefore, snow is white.

The premises of this "argument" are both true, and the conclusion is true as well. Yet it is obviously a bad argument, because it is not valid—the conclusion does not follow from the premises. The point is that *when we ask whether an argument is valid, we are not asking whether the premises actually are true, or whether the conclusion actually is true. We are only asking whether,* if *the premises were true, the conclusion would really follow from them.*

   I have used these trivial examples because they permit us to make the essential logical points clearly and uncontroversially. But these points are applicable to the analysis of any argument, trivial or not. To illustrate, let us consider how these points can be used to analyze a more important and controversial issue. We will look at the arguments for Moral Skepticism in some detail.

MORAL SKEPTICISM

Moral Skepticism is the idea that *there is no such thing as objective moral truth.* It is not merely the idea that we cannot *know* the truth about right and wrong. It is the more radical idea that, where ethics is concerned, "truth" does not exist. The essential point may be put in several different ways. It may be said that:

Morality is subjective; it is a matter of how we feel about things, not a matter of how things *are.*
Morality is only a matter of opinion, and one person's opinion is just as good as another's.
Values exist only in our minds, not in the world outside us.

However the point is put, the underlying thought is the same: The idea of "objective moral truth" is only a fiction; in reality, there is no such thing.

   We want to know whether Moral Skepticism is correct. Is the idea of moral "truth" only an illusion? What arguments can be given in favor of this idea? In order to determine whether it is correct, we need to ask

what arguments can be given for it and whether those arguments are sound.

The Cultural Differences Argument

One argument for Moral Skepticism might be based on the observation that in different cultures people have different ideas concerning right and wrong. For example, in traditional Eskimo society, infanticide was thought to be morally acceptable—if a family already had too many children, a new baby might be left to die in the snow. In our own society, however, this would be considered wrong. There are many other examples of the same kind. Different cultures have different moral codes.

Reflecting on such facts, many people have concluded that there is no such thing as objective right and wrong. Thus, they advance the following argument:

(4) In some societies, such as among the Eskimos, infanticide is thought to be morally acceptable.

In other societies, such as our own, infanticide is thought to be morally vile.

Therefore, infanticide is neither objectively right nor objectively wrong; it is merely a matter of opinion that varies from culture to culture.

We may call this the "Cultural Differences Argument." This kind of argument has been tremendously influential; it has persuaded many people to be skeptical of the whole idea of "moral truth." But is it a *sound* argument? We may ask two questions about it: First, are the premises true, and second, does the conclusion really follow from them? If the answer to either question is no, then the argument must be rejected. In this case, the premises seem to be correct—there have been many cultures in which infanticide was accepted. Therefore, our attention must focus on the second matter: Is the argument *valid*?

To figure this out, we may begin by noting that the premises concern *what people believe*. In some societies, people think infanticide is all right. In others, people believe it is immoral. The conclusion, however, concerns not what people believe, but whether infanticide *really is* immoral. The problem is that this sort of conclusion does not follow from those sorts of premises. It does not follow, from the mere fact that

people have different beliefs about something, that there is no "truth" in the matter. Therefore, the Cultural Differences Argument is not valid.

To make this point clearer, consider this analogous argument:

(5) In some societies, the world is thought to be flat.
In other societies, the world is thought to be round.
Therefore, objectively speaking, the world is neither flat nor round. It is merely a matter of opinion that varies from culture to culture.

Clearly, *this* argument is not valid. We cannot conclude that the world is shapeless simply because not everyone agrees what shape it has. But exactly the same can be said about the Cultural Differences Argument: We cannot validly move from premises about what people believe to a conclusion about what is so, because people—even whole societies— may be wrong. The world has a definite shape, and those who think it is flat are mistaken. Similarly, infanticide might be objectively wrong (or not wrong), and those who think differently might be mistaken. Therefore, the Cultural Differences Argument is not valid, and so it provides no legitimate support for the idea that "moral truth" is only an illusion.

There are two common reactions to this analysis. These reactions illustrate traps that people often fall into.

1. The first reaction goes like this: Many people find the conclusion of the Cultural Differences Argument very appealing. This makes it hard for them to believe that the argument is invalid—when the argument is shown to be fallacious, they tend to respond: "But right and wrong really *are* only matters of opinion!" They make the mistake of thinking that if we reject an argument, we are somehow impugning the truth of its conclusion. But that is not so. Remember example (3) above; it illustrates how an argument may have a true conclusion and still be a bad argument. If an argument is unsound, then it fails to provide any reason for thinking the conclusion is true. The conclusion may still be true—that remains an open question— but the point is just that the unsound argument gives it no support.

2. One may object that it is unfair to compare morality with an obviously objective matter like the shape of the earth, because we can prove what shape the earth has by scientific methods. Therefore, we know that the flat-earthers are simply wrong. But morality is different. There is no way to prove that a moral opinion is true or false.

This objection misses the point. The Cultural Differences Argument tries to derive the skeptical conclusion about morality *from a certain set of facts,* namely, the facts about cultural disagreements. This objection suggests that the conclusion might be derived from a *different* set of facts, namely, facts about what is and what is not provable. It suggests, in effect, a different argument, which might be formulated like this:

> (**6**) If infanticide (or anything else, for that matter) is objectively right or wrong, then it should be possible to *prove* that it is right or wrong.
>
> But it is not possible to prove that infanticide is right or wrong.
>
> Therefore, infanticide is neither objectively right nor objectively wrong. It is merely a matter of opinion that varies from culture to culture.

This argument is fundamentally different from the Cultural Differences Argument, even though the two arguments have the same conclusion. They are different because they appeal to different considerations in trying to prove that conclusion—in other words, they have different premises. Therefore, the question of whether argument (6) is sound is separate from the question of whether the Cultural Differences Argument is sound. The Cultural Differences Argument is not valid (and, therefore, is not sound), for the reason given above.

We should emphasize the importance of *keeping arguments separate.* It is easy to slide from one argument to another without realizing what one is doing. It is easy to think that if moral judgments are "unprovable," then the Cultural Differences Argument is strengthened. But it is not. Argument (6) merely introduces a different set of issues. It is important to pin down an argument and evaluate *it* as carefully as possible, before moving on to different considerations.

The Provability Argument

Now let us consider in more detail the question of whether it is possible to prove a moral judgment true or false. The following argument, which we might call the "Provability Argument," is a more general form of argument (6):

> (**7**) If there were any such thing as objective truth in ethics, we should be able to prove that some moral opinions are true and others false.

But, in fact, we cannot prove which moral opinions are true and which are false.

Therefore, there is no such thing as objective truth in ethics.

Once again, we have an argument with a certain superficial appeal. But are the premises true? And does the conclusion really follow from them? It seems that the conclusion does follow. Therefore, the crucial question will be whether the premises are in fact true.

The general claim that moral judgments can't be proven *sounds* right: Anyone who has ever argued about a matter like abortion knows how frustrating it can be to try to "prove" that one's point of view is correct. However, if we inspect this claim more closely, it turns out to be questionable.

Suppose we consider a matter that is simpler than abortion. A student says that a test given by a teacher was unfair. This is clearly a moral judgment—fairness is a basic moral value. Can the student prove the test was unfair? She might point out that the test was so long that not even the best students could complete it in the time allowed (and the test was to be graded on the assumption that it should be completed). Moreover, the test covered trivial matters while ignoring matters the teacher had stressed as important. And finally, the test included questions about some matters that were not covered in either the assigned readings or the class discussions.

Suppose all this is true. And further suppose that the teacher, when asked to explain, has no defense to offer. (In fact, the teacher seems confused about the whole thing and doesn't seem to have any idea what he was doing.) Now, hasn't the student proved the test was unfair? What more in the way of proof could we possibly want?

It is easy to think of other examples that make the same point:

Jones is a bad man. To prove this, one might point out that Jones is a habitual liar; he manipulates people; he cheats when he thinks he can get away with it; he is cruel to other people; and so on.

Dr. Smith is irresponsible. She bases her diagnoses on superficial considerations; she drinks Budweiser before performing delicate surgery; she refuses to listen to other doctors' advice; and so on.

A certain used-car salesman is unethical. He conceals defects in his cars; he takes advantage of poor people by pressuring them into paying high prices for cars he knows are defective; he runs false advertisements on the web; and so on.

The point is that we can, and often do, back up our ethical judgments with good reasons. Thus, it does not seem right to say that they are all unprovable, as though they were nothing more than "mere opinions." If a person has good reasons for his judgments, then he is not *merely* giving "his opinion." On the contrary, he may be making a judgment which any reasonable person would have to agree with.

If we can sometimes give good reasons for our moral judgments, what accounts for the persistent impression that they are "unprovable"? There are two reasons why the Provability Argument appears better than it is.

First, there is a tendency to focus attention only on the most difficult moral issues. The question of abortion, for example, is an enormously difficult and complicated matter. If we think only of questions like *this,* it is easy to believe that "proof" in ethics is impossible. The same could be said of the sciences. There are many complicated matters that physicists cannot agree on; if we focused our attention entirely on *them,* we might conclude that there is no "proof" in physics. But, of course, there are many simpler matters in physics that can be proven and about which all competent physicists agree. Similarly, in ethics, there are many matters far simpler than abortion, about which all reasonable people must agree.

Second, it is easy to confuse two matters that are really very different:

1. Proving an opinion to be correct
2. Persuading someone to accept your proof

Suppose you are discussing a moral issue with a friend. You have perfectly cogent reasons in support of your position, while he has no good reasons on his side. Still, he refuses to accept your logic and continues to insist that he is right. This is a common, if frustrating, experience. You may be tempted to conclude that it is impossible to prove you are right. But this would be a mistake. Your proof may be impeccable; the trouble may be that your friend is being stubborn. (Of course, that is not the *only* possible explanation of what is going on, but it is one possible explanation.) The same thing can happen in any sort of discussion. You may be arguing about intelligent design versus evolution, and the other person may be unreasonable. But that does not necessarily mean that something is wrong with your arguments. Something may be wrong with the other person.

CONCLUSION

We have examined two of the most important arguments in support of Moral Skepticism and seen that these arguments are no good. Moral Skepticism might still turn out to be true, but if so, then other, better arguments will have to be found. Provisionally, at least, we have to conclude that Moral Skepticism is not nearly as plausible as we might have thought.

The purpose of this exercise, however, was to illustrate the process of evaluating philosophical arguments. We may summarize the main points like this:

1. Arguments are offered to provide support for a theory or idea; a philosophical theory may be regarded as acceptable only if there are sound arguments in its favor.
2. An argument is sound only if its premises are true and the conclusion follows logically from them.
   (a) A conclusion "follows from" the premises just in case the following is so: *If* the premises were true, then the conclusion would have to be true also. (An alternative way of saying the same thing: A conclusion follows from the premises just in case it is impossible for the premises to be true and the conclusion false at the same time.)
   (b) A conclusion can follow from premises even if those premises are in fact false.
   (c) A conclusion can be true and yet not follow from a given set of premises.
3. Therefore, in evaluating an argument, we ask two *separate* questions: Are the premises true? and Does the conclusion follow from them?
4. It is important to avoid two common mistakes. We should be careful to keep arguments separate, and not slide from one to the other, thereby confusing different issues. And, we should not think an argument stronger than it is simply because we happen to agree with its conclusion. Moreover, we should remember that, if an argument is unsound, that does not mean the conclusion must be false—it only means that *this* argument does nothing to show that the conclusion is true.

IS MORALITY OBJECTIVE?

Chapter 3

The Subjectivity of Values

J. L. Mackie

Everyone agrees that ethics is subjective in the sense that people have their own personal moral beliefs, and that those beliefs differ. But in this selection, John L. Mackie contends that ethics is subjective in a much more radical sense—that, really, there is no right or wrong.

Consider this analogy: The earth is round (or spherical), and not flat, because there is this thing—the earth—that has the property of being round. But when someone says abortion is wrong, according to Mackie, there can be no property of wrongness that inheres in abortion in analogy to how roundness adheres in the earth; there is just the feeling, or belief, that abortion is wrong.

In Mackie's view, the belief that Abraham Lincoln is morally better than John Wilkes Booth is similar to the cartoon character Homer Simpson's belief that syrup is better than jelly. Homer might like syrup more than jelly, but he's not "right" in feeling that way. Similarly, Mackie thinks, you might approve of the president who freed the slaves more than you approve of the man who shot him, but you are not "right" in feeling that way. It's just how you feel.

John L. Mackie (1917–1981) was born in Australia and taught at the University of Oxford. This selection is from his book *Ethics: Inventing Right and Wrong* (1977).

MORAL SCEPTICISM

There are no objective values. This is a bald statement of the thesis of this chapter, but before arguing for it I shall try to clarify and restrict it in ways that may meet some objections and prevent some misunderstanding.

The statement of this thesis is liable to provoke one of three very different reactions. Some will think it not merely false but pernicious; they will see it as a threat to morality and to everything else that is worthwhile, and they will find the presenting of such a thesis in what purports to be a book on ethics paradoxical or even outrageous. Others will regard it as a trivial truth, almost too obvious to be worth mentioning, and certainly too plain to be worth much argument. Others again will say that it is meaningless or empty, that no real issue is raised by the question whether values are or are not part of the fabric of the world. But, precisely because there can be these three different reactions, much more needs to be said.

The claim that values are not objective, are not part of the fabric of the world, is meant to include not only moral goodness, which might be most naturally equated with moral value, but also other things that could be more loosely called moral values or disvalues—rightness and wrongness, duty, obligation, an action's being rotten and contemptible, and so on. It also includes non-moral values, notably aesthetic ones, beauty and various kinds of artistic merit. I shall not discuss these explicitly, but clearly much the same considerations apply to aesthetic and to moral values, and there would be at least some initial implausibility in a view that gave the one a different status from the other.

Since it is with moral values that I am primarily concerned, the view I am adopting may be called moral scepticism. But this name is likely to be misunderstood: "moral scepticism" might also be used as a name for either of two first order views, or perhaps for an incoherent mixture of the two. A moral sceptic might be the sort of person who says "All this talk of morality is tripe," who rejects morality and will take no notice of it. Such a person may be literally rejecting all moral judgements; he is more likely to be making moral judgements of his own, expressing a positive moral condemnation of all that conventionally passes for morality; or he may be confusing these two logically incompatible

views, and saying that he rejects all morality, while he is in fact reject-ing only a particular morality that is current in the society in which he has grown up. But I am not at present concerned with the merits or faults of such a position. These are first order moral views, positive or negative: the person who adopts either of them is taking a certain prac-tical, normative, stand. By contrast, what I am discussing is a second order view, a view about the status of moral values and the nature of moral valuing, about where and how they fit into the world. These first and second order views are not merely distinct but completely indepen-dent: one could be a second order moral sceptic without being a first order one, or again the other way round. A man could hold strong moral views, and indeed ones whose content was thoroughly conventional, while believing that they were simply attitudes and policies with regard to conduct that he and other people held. Conversely, a man could reject all established morality while believing it to be an objective truth that it was evil or corrupt.

With another sort of misunderstanding moral scepticism would seem not so much pernicious as absurd. How could anyone deny that there is a difference between a kind action and a cruel one, or that a coward and a brave man behave differently in the face of danger? Of course, this is undeniable; but it is not to the point. The kinds of behaviour to which moral values and disvalues are ascribed are indeed part of the furniture of the world, and so are the natural, descriptive, differences between them; but not, perhaps, their differences in value. It is a hard fact that cruel actions differ from kind ones, and hence that we can learn, as in fact we all do, to distinguish them fairly well in practice, and to use the words "cruel" and "kind" with fairly clear descriptive meanings; but is it an equally hard fact that actions which are cruel in such a descrip-tive sense are to be condemned? The present issue is with regard to the objectivity specifically of value, not with regard to the objectivity of those natural, factual, differences on the basis of which differing values are assigned. . . .

STANDARDS OF EVALUATION

One way of stating the thesis that there are no objective values is to say that value statements cannot be either true or false. But this formulation,

too, lends itself to misinterpretation. For there are certain kinds of value statements which undoubtedly can be true or false, even if, in the sense I intend, there are no objective values. Evaluations of many sorts are commonly made in relation to agreed and assumed standards. The classing of wool, the grading of apples, the awarding of prizes at sheepdog trials, flower shows, skating and diving championships, and even the marking of examination papers are carried out in relation to standards of quality or merit which are peculiar to each particular subject-matter or type of contest, which may be explicitly laid down but which, even if they are nowhere explicitly stated, are fairly well understood and agreed upon by those who are recognized as judges or experts in each particular field. Given any sufficiently determinate standards, it will be an objective issue, a matter of truth and falsehood, how well any particular specimen measures up to those standards. Comparative judgements in particular will be capable of truth and falsehood: it will be a factual question whether this sheepdog has performed better than that one.

The subjectivist about values, then, is not denying that there can be objective evaluations relative to standards, and these are as possible in the aesthetic and moral fields as in any of those just mentioned. More than this, there is an objective distinction which applies in many such fields, and yet would itself be regarded as a peculiarly moral one: the distinction between justice and injustice. In one important sense of the word it is a paradigm case of injustice if a court declares someone to be guilty of an offence of which it knows him to be innocent. More generally, a finding is unjust if it is at variance with what the relevant law and the facts together require, and particularly if it is known by the court to be so. More generally still, any award of marks, prizes, or the like is unjust if it is at variance with the agreed standards for the contest in question: if one diver's performance in fact measures up better to the accepted standards for diving than another's, it will be unjust if the latter is awarded higher marks or the prize. In this way the justice or injustice of decisions relative to standards can be a thoroughly objective matter, though there may still be a subjective element in the interpretation or application of standards. But the statement that a certain decision is thus just or unjust will not be objectively prescriptive: in so far as it can be simply true it leaves open the question whether there is any objective

requirement to do what is just and to refrain from what is unjust, and equally leaves open the practical decision to act in either way.

Recognizing the objectivity of justice in relation to standards, and of evaluative judgements relative to standards, then, merely shifts the question of the objectivity of values back to the standards themselves. The subjectivist may try to make his point by insisting that there is no objective validity about the choice of standards. Yet he would clearly be wrong if he said that the choice of even the most basic standards in any field was completely arbitrary. The standards used in sheepdog trials clearly bear some relation to the work that sheepdogs are kept to do, the standards for grading apples bear some relation to what people generally want in or like about apples, and so on. On the other hand, standards are not as a rule strictly validated by such purposes. The appropriateness of standards is neither fully determinate nor totally indeterminate in relation to independently specifiable aims or desires. But however determinate it is, the objective appropriateness of standards in relation to aims or desires is no more of a threat to the denial of objective values than is the objectivity of evaluation relative to standards. In fact it is logically no different from the objectivity of goodness relative to desires. Something may be called good simply in so far as it satisfies or is such as to satisfy a certain desire; but the objectivity of such relations of satisfaction does not constitute in our sense an objective value.

HYPOTHETICAL AND CATEGORICAL
IMPERATIVES

We may make this issue clearer by referring to Kant's distinction between hypothetical and categorical imperatives, though what he called imperatives are more naturally expressed as "ought" statements than in the imperative mood. "If you want *X*, do *Y*" (or "You ought to do *Y*") will be a hypothetical imperative if it is based on the supposed fact that *Y* is, in the circumstances, the only (or the best) available means to *X*, that is, on a causal relation between *Y* and *X*. The reason for doing *Y* lies in its causal connection with the desired end, *X*; the oughtness is contingent upon the desire. But "You ought to do *Y*" will be a categorical imperative if you ought to do *Y* irrespective of any such

desire for any end to which *Y* would contribute, if the oughtness is not thus contingent upon any desire.

A categorical imperative, then, would express a reason for acting which was unconditional in the sense of not being contingent upon any present desire of the agent to whose satisfaction the recommended action would contribute as a means—or more directly: "You ought to dance," if the implied reason is just that you want to dance or like dancing, is still a hypothetical imperative. Now Kant himself held that moral judgements are categorical imperatives, or perhaps are all applications of one categorical imperative, and it can plausibly be maintained at least that many moral judgements contain a categorically imperative element. So far as ethics is concerned, my thesis that there are no objective values is specifically the denial that any such categorically imperative element is objectively valid. The objective values which I am denying would be action-directing absolutely, not contingently (in the way indicated) upon the agent's desires and inclinations.

Another way of trying to clarify this issue is to refer to moral reasoning or moral arguments. In practice, of course, such reasoning is seldom fully explicit: but let us suppose that we could make explicit the reasoning that supports some evaluative conclusion, where this conclusion has some action-guiding force that is not contingent upon desires or purposes or chosen ends. Then what I am saying is that somewhere in the input to this argument—perhaps in one or more of the premisses, perhaps in some part of the form of the argument—there will be something which cannot be objectively validated—some premiss which is not capable of being simply true, or some form of argument which is not valid as a matter of general logic, whose authority or cogency is not objective, but is constituted by our choosing or deciding to think in a certain way.

THE CLAIM TO OBJECTIVITY

If I have succeeded in specifying precisely enough the moral values whose objectivity I am denying, my thesis may now seem to be trivially true. Of course, some will say, valuing, preferring, choosing, recommending, rejecting, condemning, and so on, are human activities, and there is no need to look for values that are prior to and logically

independent of all such activities. There may be widespread agreement in valuing, and particular value-judgements are not in general arbitrary or isolated: they typically cohere with others, or can be criticized if they do not, reasons can be given for them, and so on: but if all that the subjectivist is maintaining is that desires, ends, purposes, and the like figure somewhere in the system of reasons, and that no ends or purposes are objective as opposed to being merely intersubjective, then this may be conceded without much fuss.

But I do not think that this should be conceded so easily. As I have said, the main tradition of European moral philosophy includes the contrary claim, that there are objective values of just the sort I have denied. . . . Kant in particular holds that the categorical imperative is not only categorical and imperative but objectively so: though a rational being gives the moral law to himself, the law that he thus makes is determinate and necessary. Aristotle begins the *Nicomachean Ethics* by saying that the good is that at which all things aim, and that ethics is part of a science which he calls "politics," whose goal is not knowledge but practice; yet he does not doubt that there can be *knowledge* of what is the good for man, nor, once he has identified this as well-being or happiness, *eudaimonia*, that it can be known, rationally determined, in what happiness consists; and it is plain that he thinks that this happiness is intrinsically desirable, not good simply because it is desired. . . . Even the sentimentalist Hutcheson defines moral goodness as "some quality apprehended in actions, which procures approbation, . . ." while saying that the moral sense by which we perceive virtue and vice has been given to us (by the Author of nature) to direct our actions. Hume indeed was on the other side, but he is still a witness to the dominance of the objectivist tradition, since he claims that when we "see that the distinction of vice and virtue is not founded merely on the relations of objects, nor is perceiv'd by reason," this "wou'd subvert all the vulgar systems of morality." . . .

The prevalence of this tendency to objectify values—and not only moral ones—is confirmed by a pattern of thinking that we find in existentialists and those influenced by them. The denial of objective values can carry with it an extreme emotional reaction, a feeling that nothing matters at all, that life has lost its purpose. Of course this does not follow; the lack of objective values is not a good reason for abandoning subjective concern or for ceasing to want anything. But

the abandonment of a belief in objective values can cause, at least temporarily, a decay of subjective concern and sense of purpose. That it does so is evidence that the people in whom this reaction occurs have been tending to objectify their concerns and purposes, have been giving them a fictitious external authority. A claim to objectivity has been so strongly associated with their subjective concerns and purposes that the collapse of the former seems to undermine the latter as well.

This view, that conceptual analysis would reveal a claim to objectivity, is sometimes dramatically confirmed by philosophers who are officially on the other side. Bertrand Russell, for example, says that "ethical propositions should be expressed [as desires]"; he defends himself effectively against the charge of inconsistency in both holding ultimate ethical valuations to be subjective and expressing emphatic opinions on ethical questions. Yet at the end he admits:

> Certainly there *seems* to be something more. Suppose, for example, that someone were to advocate the introduction of bullfighting in this country. In opposing the proposal, I should *feel*, not only that I was expressing my desires, but that my desires in the matter are *right*, whatever that may mean. As a matter of argument, I can, I think, show that I am not guilty of any logical inconsistency in holding to the above interpretation of ethics and at the same time expressing strong ethical preferences. But in feeling I am not satisfied.

But he concludes, reasonably enough, with the remark: "I can only say that, while my own opinions as to ethics do not satisfy me, other people's satisfy me still less."

I conclude, then, that ordinary moral judgements include a claim to objectivity, an assumption that there are objective values in just the sense in which I am concerned to deny this. And I do not think it is going too far to say that this assumption has been incorporated in the basic, conventional, meanings of moral terms. Any analysis of the meanings of moral terms which omits this claim to objective, intrinsic, prescriptivity is to that extent incomplete. . . .

If second order ethics were confined, then, to linguistic and conceptual analysis, it ought to conclude that moral values at least are objective: that they are so is part of what our ordinary moral statements mean: the traditional moral concepts of the ordinary man as well as of the main line of Western philosophers are concepts of objective value.

But it is precisely for this reason that linguistic and conceptual analysis is not enough. The claim to objectivity, however ingrained in our language and thought, is not self-validating. It can and should be questioned. But the denial of objective values will have to be put forward not as the result of an analytic approach, but as an "error theory," a theory that although most people in making moral judgements implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false. It is this that makes the name "moral scepticism" appropriate.

But since this is an error theory, since it goes against assumptions ingrained in our thought and built into some of the ways in which language is used, since it conflicts with what is sometimes called common sense, it needs very solid support. It is not something we can accept lightly or casually and then quietly pass on. If we are to adopt this view, we must argue explicitly for it. Traditionally it has been supported by arguments of two main kinds, which I shall call the argument from relativity and the argument from queerness, but these can, as I shall show, be supplemented in several ways.

THE ARGUMENT FROM RELATIVITY

The argument from relativity has as its premiss the well-known variation in moral codes from one society to another and from one period to another, and also the differences in moral beliefs between different groups and classes within a complex community. Such variation is in itself merely a truth of descriptive morality, a fact of anthropology which entails neither first order nor second order ethical views. Yet it may indirectly support second order subjectivism: radical differences between first order moral judgements make it difficult to treat those judgements as apprehensions of objective truths. But it is not the mere occurrence of disagreements that tells against the objectivity of values. Disagreement on questions in history or biology or cosmology does not show that there are no objective issues in these fields for investigators to disagree about. But such scientific disagreement results from speculative inferences or explanatory hypotheses based on inadequate evidence, and it is hardly plausible to interpret moral disagreement in the same way. Disagreement about moral codes seems to reflect people's

adherence to and participation in different ways of life. The causal con-
nection seems to be mainly that way round: it is that people approve of
monogamy because they participate in a monogamous way of life rather
than that they participate in a monogamous way of life because they
approve of monogamy. Of course, the standards may be an idealiza-
tion of the way of life from which they arise: the monogamy in which
people participate may be less complete, less rigid, than that of which
it leads them to approve. This is not to say that moral judgements are
purely conventional. Of course there have been and are moral heretics
and moral reformers, people who have turned against the established
rules and practices of their own communities for moral reasons, and
often for moral reasons that we would endorse. But this can usually be
understood as the extension, in ways which, though new and uncon-
ventional, seemed to them to be required for consistency, of rules to
which they already adhered as arising out of an existing way of life.
In short, the argument from relativity has some force simply because
the actual variations in the moral codes are more readily explained by
the hypothesis that they reflect ways of life than by the hypothesis that
they express perceptions, most of them seriously inadequate and badly
distorted, of objective values.

But there is a well-known counter to this argument from relativ-
ity; namely to say that the items for which objective validity is in the
first place to be claimed are not specific moral rules or codes but very
general basic principles which are recognized at least implicitly to
some extent in all society—such principles provide the foundations of
what Sidgwick has called different methods of ethics: the principle of
universalizability, perhaps, or the rule that one ought to conform to the
specific rules of any way of life in which one takes part, from which
one profits, and on which one relies, or some utilitarian principle of
doing what tends, or seems likely, to promote the general happiness.
It is easy to show that such general principles, married with differing
concrete circumstances, different existing social patterns or different
preferences, will beget different specific moral rules; and there is some
plausibility in the claim that the specific rules thus generated will vary
from community to community or from group to group in close agree-
ment with the actual variations in accepted codes.

The argument from relativity can be only partly countered in this
way. To take this line the moral objectivist has to say that it is only in