
Ethics and Objectivity

There are no objective values.

—J. L. MACKIE, *ETHICS: INVENTING RIGHT AND WRONG* (1977)

11.1. Thrasymachus's Challenge

Thrasymachus has the misfortune of being remembered through the eyes of someone who despised him. He was a Sophist, one of the professional teachers who flourished in Athens during the time of Socrates (ca. 470–399 B.C.). We know about Socrates through the writings of his student Plato. Plato tells us that the Sophists, unlike Socrates, charged a fee for their instruction, and he is quick to insinuate that the Sophists liked money more than truth. Plato is especially hard on Thrasymachus, who is introduced in the *Republic* like this:

While we had been talking [says Socrates] Thrasymachus had often tried to interrupt, but had been prevented by those sitting near him, who wanted to hear the argument concluded; but when we paused, he was no longer able to contain himself and gathered himself together and sprang on us like a wild beast, as if he wanted to tear us in pieces. Polemarchus and I were scared stiff, as Thrasymachus burst out and said, “What is all this nonsense, Socrates?”

The “nonsense” was a discussion of justice. Thrasymachus was impatient because Socrates and his friends were assuming that justice is something real and important. But according to Thrasymachus, people believe in right and wrong only because they are taught to obey the rules of their society. Those rules, however, are merely human inventions. Thrasymachus added that the rules of a society will protect the interests of the society's

most powerful members. So, when ordinary people think they must “do the right thing,” they are just being played.

Throughout history, there have always been groups of people who, like Thrasymachus, believe that ethics is just a matter of opinion; and there have always been groups of people who, like Socrates, believe that ethics has an objective basis. But with the rise of modern science, skepticism about ethics became even more attractive. Modern science sees the world as a cold, indifferent place that cares nothing for us or our projects; the universe is a realm of facts that know nothing of right or wrong. As David Hume (1711–1776) put it, “The life of a man is of no greater importance to the universe than that of an oyster.” Thus, it seems natural to conclude that ethics can be nothing but a human creation.

Thrasymachus challenged Socrates to prove that ethics has an objective basis. One way of meeting Thrasymachus’s challenge might be to introduce religious notions. If the universe was created by God, according to a divine plan, and if God issues commands about how we should live, we might find in this an objective basis for our judgments of right and wrong. But suppose we set this possibility aside. Is there any way to defend the objectivity of ethics without invoking religion? We will see that there is a way. The arguments for ethical skepticism are not as powerful as they appear.

11.2. Is Ethics Just a Matter of Social Conventions?

The idea that ethics is nothing more than social convention has always appealed to educated people. Different cultures have different moral codes, it is said, and it is merely naïve to think that there is one universal standard that applies in all places and times. Examples of differences are easy to come by. In Islamic countries, men may have more than one wife. In medieval Europe, lending money for interest was considered a sin. In northern Greenland, old people were sometimes left to die in the snow. Considering such examples, anthropologists have long agreed with Herodotus, the ancient historian who said, “Custom is king o’er all.”

Today, the idea of morality as a social product is attractive for an additional reason. Given America’s economic power and military might, it is especially important for Americans to respect

and appreciate the differences between cultures. In particular, it is said, we must avoid the arrogant assumption that our ways are “right” and that other ways are “wrong.” This means, in part, that we should refrain from making moral judgments about other cultures. We should adopt a policy of live and let live.

On the surface, this attitude seems enlightened. Tolerance is important, and many cultural practices are nothing more than social customs—standards of dress, the details of household arrangements, the methods of greeting, and so on. But fundamental matters of justice are different. When we consider such examples as slavery, racism, and the abuse of women, it no longer seems enlightened to shrug and say, “They have their customs, and we have ours.” Consider two examples.

In a Pakistani village, a 12-year-old boy was accused of being romantically involved with a 22-year-old woman of a higher social class. He denied it, but the tribal elders did not believe him. As punishment, they decreed that the boy’s teenage sister—who had done nothing wrong—should be publicly raped. Her name is Mukhtar Mai. Four men carried out the sentence while the village watched. Observers said there was nothing unusual in this, but with so many foreigners in the region, the incident was noticed and reported in *Newsweek*. This was in 2002.

In the same year, in northern Nigeria, a religious court sentenced an unwed mother named Amina Lawal to be stoned to death for having had sex out of wedlock. When the verdict was read, the crowd in the courtroom cheered, and the judge said that the sentence should be carried out as soon as Lawal’s baby was old enough not to need breast-feeding. Lawal identified the father, but he denied the accusation, and no charges were brought against him. This was only one in a series of such sentences imposed in northern Nigeria. Responding to international pressure, the Nigerian government announced that it would not enforce the sentence against Lawal, and in 2004 she was set free.

The rape of Mukhtar Mai was regarded as a matter of tribal honor. Her brother was allegedly romancing a woman from a different tribe, and the elders of her tribe demanded justice. The stonings in Nigeria, on the other hand, are the application of the Islamic law of Sharia, which has been adopted by 12 of Nigeria’s 36 states. Our instincts are to condemn both punishments. But are we *justified* in condemning them? Two thoughts stand in the way of saying we are.

First, there is the idea that *we should respect the differences between cultures*. People in other cultures have a right to follow their own traditions, it is said, even if outsiders like us disapprove. After all, their traditions may have a purpose that we don't understand. Moreover, it is said, their values might be different from ours. Should thoughts like these persuade us to hold our tongues when we hear about a raping or a stoning on the other side of the world?

They should not. Respecting a culture does not mean tolerating everything in it. You might think that a culture has a wonderful history and has produced great art. You might think that its leading figures are admirable and that your own culture has much to learn from them. Yet, despite all this, you need not regard the culture as perfect. You might think it contains elements that are awful. Most of us take just this attitude toward our own society—if you are an American, you probably think that America is a great country but that some aspects of our culture are appalling. Why should you not think the same about Pakistan or Nigeria? If you did, you would be agreeing with many Pakistanis and Nigerians.

Moreover, it is a mistake to think of the world as a collection of discrete, unified cultures that exist in isolation from one another. Cultures overlap and interact. In the United States, there are cultural differences between Irish Catholics, Italian Americans, Southern Baptists, African Americans in Los Angeles, African Americans in Mississippi, and Hasidic Jews in Brooklyn. Coal miners in West Virginia are quite different, culturally, from stockbrokers in New York City. In some ways we think that “live and let live” is the best policy, but no one takes this to mean that you should have no opinion about what happens in another part of the country.

This also raises the question of who speaks for a culture. Is it the priests? The politicians? The women? The slaves? Opinions within a society are rarely uniform. If we say, for example, that slavery was approved of in ancient Greece, we are referring to the opinions of the slave-owners. The slaves themselves might have had a different opinion. Or consider again the public raping of Mukhtar Mai. When this happened, the Pakistani government took action against the tribal leaders who had ordered it. Which group—the local leaders or the national government—sets the standards that we must respect? There are no clear-cut answers to these questions. Thus, it is often unhelpful to say, “We must respect the values of the culture.”

Finally, we should notice a purely logical point. Some people think that ethical relativism *follows from* the fact that cultures have different standards. That is, they endorse this argument:

- (1) Different cultures have different moral codes.
- (2) Therefore, there is no such thing as objective right and wrong. Where ethics is concerned, the standards of the different societies are all that exist.

But this is a mistake. It does not follow from the fact that people disagree about something, that there is no truth about it. For example, cultures may disagree about the Milky Way—some think it's a galaxy, others think it's a river in the sky—but it does not follow from this that there is no objective truth about the Milky Way. The same goes for ethics. When cultures employ different customs, some of the customs may be better than others. It is easy to overlook this if we only think of trivial examples, such as the standards of dress at a wedding or funeral. Those may indeed be nothing but matters of local custom. But it does not follow that *all* practices are merely matters of local custom. Fundamental matters of justice might be different.

Thus, we needn't refrain from morally assessing the customs of other societies. We can be tolerant and respectful, yet think that other cultures are flawed.

There is, however, a second reason why being judgmental may seem inappropriate: all standards may seem to be culture-relative. If we say that the rape of Mukhtar Mai was wrong, we are using *our* standards to judge *their* practices. From our point of view, the rape was wrong, but why is our point of view correct? We can say that the tribal leaders are wrong, but they can equally well say that we are wrong. How can we get beyond this mutual finger-pointing?

This second argument can be spelled out as follows:

- (1) For our criticisms of other cultures to be justified, they must appeal to standards that are not simply derived from our own culture.
- (2) But there are no such culture-neutral moral standards. All standards are relative to some society or other.
- (3) Therefore, our criticisms of other cultures are unjustified.

Is this correct? It sounds plausible, but in fact there *is* a culture-neutral standard of right and wrong. The reason we object to the rape and the stoning is not that they are “contrary to American values.” Nor is our objection that these practices are somehow bad for *us*. We object because of the harm done to Mukhtar Mai and Amina Lawal. Our culture-neutral standard is *whether the social practice in question is beneficial or harmful to the people who are affected by it*. Good social practices benefit people; bad social practices harm people.

This standard is culture-neutral in every relevant sense. First, it does not play favorites between cultures. It may be applied equally to all societies, including our own. Second, the source of the principle does not lie within one particular culture. Rather, every culture values the welfare of its people. It is a value that must be embraced, at least to some extent, by every culture, if the culture is to exist. And so, the suggestion that a social practice is harmful can never be dismissed as an alien standard “brought in from the outside.”

11.3. Ethics and Science

We think of science as the most objective human enterprise. Scientists know how to get at the truth. The study of ethics, however, seems unlike science. So how can ethics be objective? Let’s consider three arguments along these lines.

The Argument from Disagreement. It is troubling that ethical disagreement seems so widespread and persistent. If ethics were a matter of objective truth, shouldn’t we expect more consensus? Yet it seems that, in matters of ethics, people disagree about everything. They argue over abortion, capital punishment, gun control, euthanasia, the environment, and the moral status of animals. They disagree about sex, drug use, and whether we have a duty to help needy children in foreign countries. In science, however, there is widespread agreement on all essential points. The natural conclusion is that ethics, unlike science, is a mere matter of opinion. We may summarize the argument like this:

- (1) In ethics, unlike in science, there is widespread and persistent disagreement.

- (2) The best explanation of this fact is that there is no objective truth in ethics.
- (3) Therefore, we may conclude, at least tentatively, that there is no objective truth in ethics.

Is this argument sound? We may begin by observing that ethics is more like science than people think. There is a tremendous amount of agreement in ethics. All thoughtful people agree that murder, rape, theft, child abuse, blackmail, kidnapping, and racism are wrong. Everyone agrees that we should tell the truth, help our friends, keep our promises, and love our children. If it is said that some people do not agree—racists and thieves, for example—it can be replied that some people disagree with the findings of science—flat-earthers and psychics, for example. The situation in ethics is the same as in science: The vast majority of people agree, while some dissenters are ignored, for good reason. In fact, there may be more dissenters in science than in ethics, if you count religious fundamentalists who reject Darwin's theory of evolution.

In ethics, then, there is massive agreement about simple matters. But there is also disagreement about abortion, capital punishment, and the other issues mentioned above. What should we make of this? We might notice, first, that from a social point of view, most of the matters we disagree on are less important than the matters we agree on. No matter how much people care about gun control, for example, gun control is less important than murder, truth telling, or promise keeping. After all, societies can function with a variety of policies on guns. But social living would be impossible without a rule against murder. Likewise, society would be impossible if people could lie and break their promises at will. To see this, try to imagine what it would be like to live in a place where people could lie, cheat, steal, and kill with impunity. That society would collapse.

We might also notice that many of the disputed issues are *harder* than the matters we agree on. In order to take a principled stand on abortion, for example, we would have to figure out whether the fetus's life is precious. Thus, we would have to answer the question: When, in the course of human development, does human life acquire its full value? Moreover, we might have to assess the importance of potentiality: Does the fact that the fetus *might one day* become a full human person affect the fetus's moral

status *today*? All this is hard enough, but to make things worse, it is not obvious that these are even the right questions to ask. We could instead ask about the rights of the pregnant woman, or about how large the human population should be. So, it is no wonder that people disagree about the morality of abortion. In such cases, the difficulty of the issues, and not the absence of “truth,” is the best explanation of why people don’t agree.

A similar pattern of agreement and disagreement exists in science. All scientists agree about a large central core of accepted truth. Yet many others issues are disputed. Scientists disagree about the path that evolution has taken, the prospects for string theory, what is really shown by the groundbreaking experiments on infant cognition, and the relation between quantum theory and classical relativity. So, contrary to superficial impressions, there does not appear to be any fundamental difference between ethics and science. Both are characterized by broad agreement alongside some disagreement.

The Argument from Lack of Proof. While scientists may disagree about some things, they agree about how to resolve their disputes. They make observations and perform experiments in order to prove or disprove their hypotheses. This means that, in science, disagreement is only temporary. But in ethics, disagreement is endless, because no one knows how to prove or disprove anything.

This argument may be summarized as follows:

- (1) If there were objective truth in ethics, then it should be possible to prove that at least some ethical opinions are true.
- (2) But it is not possible to prove that any ethical opinion is true.
- (3) Therefore, there is no objective truth in ethics.

Is this correct? It certainly sounds plausible. Anyone who has tried to persuade someone else on an ethical matter will know how frustrating it can be. A pacifist, for example, will not be persuaded that violence is sometimes necessary, no matter what reasons are offered. Nor will a defender of violence be persuaded to change his mind, no matter how often he is told that violence only begets more violence.

If we turn to simpler examples, however, things look very different. Suppose the issue is whether a certain doctor is unethical. I say that Dr. Jones behaves shamefully, and you are surprised to hear it because you think she's a fine doctor. So I point out several things:

- Dr. Jones owns stock in a drug company, and she always prescribes that company's drugs, whether her patients need them or not.
- She won't listen to the advice of other physicians, and she becomes angry when her patients want a second opinion.
- She doesn't keep up with current medical knowledge.
- She often performs surgery while she's drunk.

Suppose all this is true. Isn't this good evidence that she is unethical? Doesn't this *prove* that she is unethical? Suppose, further, that little could be said on the other side, in her defense. Doesn't this settle the matter? What more in the way of proof could anyone want?

Other examples come easily to mind. The proof that Mr. Smith is a bad man is that he is a habitual liar who is sometimes cruel. The proof that Mr. Brown is an unethical poker player is that he cheats. The proof that Professor Adams should not have given the midterm on Tuesday is that she announced it would be given on Thursday. In each case, of course, further facts may need to be considered. But the point is that such judgments are not merely "subjective."

Ethical proofs may be different, in some ways, from scientific proofs. But that does not mean that ethical proofs are deficient. Ethical proofs consist in giving reasons to support moral conclusions. If the reasons are powerful, and if there are no good opposing considerations, then the case is made.

This may seem too quick. If ethical judgments can so obviously be proven, then why was the contrary idea so plausible in the first place? Why is it so intuitively appealing to think that there are no ethical proofs? There are at least three reasons.

First, when we think about ethics, we don't usually think about the simple matters. The fact that they are simple makes them boring, so we tend to ignore them. We are attracted instead to the harder issues, such as pacifism, abortion, and gun control. But this is what leads us astray. If we think only about the hard issues, we may naturally conclude that there are no

proofs in ethics, because no one has a knockdown argument against abortion or in favor of pacifism. We might overlook the fact that proofs are easily available on the more mundane issues.

Second, there are often good reasons on both sides of a moral issue, and this leads people to despair about reaching definite conclusions. If I say that Smith is a bad man because he often lies, you may reply that Smith sometimes works for good causes. The first fact counts against him, but the second is in his favor. This feature of moral reasoning, however, should not worry us. Moral reasoning requires taking all the facts into account and weighing them against one another. Where Smith is concerned, the right conclusion might be that he is on the whole bad, even though he has some good qualities. Or the right conclusion might be that Smith is equally good and bad. It just depends on the facts. The difficult issues are like this; there is much to be said on each side.

That simple truth is often overlooked. It is common for people on one side of a debate to deny that the other side has made any good points at all. To win the debate, they feel that they must concede nothing. This, of course, means that the debate will go on endlessly. But moral debates needn't be like that. You can acknowledge that your opponent has made good, relevant points, while maintaining that your points are even better.

Finally, it is easy to confuse *proving an opinion to be true* and *persuading someone to accept your proof*. The first is a matter of sound reasoning; the second is a matter of psychology. Someone might reject a perfectly good argument because he is stubborn, or prejudiced, or simply uninterested in finding the truth. It does not follow that the argument itself is defective. A Klansman or a neo-Nazi may not listen to a sound argument about racism, but that says something about *them*, not about the argument. And there is a more general reason why people might resist listening. Accepting a moral argument often means acknowledging that we should change our behavior. And we may not want to do that. Thus, people sometimes turn a deaf ear.

The Metaphysical Argument. There is one further argument to consider, namely, that ethics cannot be objective because “values” do not exist as part of the objective world.

If we take an inventory of the world, noting all the things that exist, we can make a very long list, mentioning rocks, rivers,

mountains, plants, and animals. We would find buildings, deserts, caves, iron, and oxygen. Looking up, we would see stars, comets, clouds, and galaxies. Of course, we could never complete such a list. Life is too short, there are too many things, and we are too ignorant. But we think we know, at least roughly, the *kinds* of things that exist. There are physical objects, made of atoms, that obey the laws of physics, chemistry, and biology; and there are conscious beings, such as ourselves, which may or may not be just another type of physical object.

But where, among all these things, are values? The answer, it seems, is *nowhere*. Values do not exist, at least not in the same way that rocks and rivers exist. Apart from human feelings and human interests, the world appears to contain no values. David Hume put it this way:

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. The vice entirely escapes you, as long as you consider the object. You never can find it, till you turn your reflection 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, not of reason.

Of course, there may be other conscious beings who also have feelings and interests—nonhuman animals, for example, and possibly the inhabitants of other planets. But they will be in the same position that we are in. They will find no values in the world around them. Only their “passions, motives, volitions, and thoughts” will give rise to values for them.

Other philosophers have taken up this theme. Friedrich Nietzsche (1844–1900) was a troubled figure who had no real home for much of his life. He moved around Europe, writing book after book, but his books were largely ignored. Eleven years before his death, he went insane. After he died, his fame grew, but he was embraced by the Nazis, who wrongly took him to support their racist causes. Later he was rescued from this misunderstanding, and today Nietzsche is seen as an original and important thinker.

Like Hume, Nietzsche denied that there are moral facts. “There are no moral phenomena at all,” he wrote, “but only

moral interpretations of phenomena.” The right way to think about ethics, he thought, is not to focus on *morality*, as though it were one unified thing, but to study *moralities*, the historically contingent systems of value that have been created by different people at different times. Nietzsche himself wrote about a morality that underlies much of Western culture, which he called “slave morality.” Slave morality, Nietzsche said, glorifies meekness, self-denial, obedience, and poverty. This outlook, he said, is unworthy of noble men, and he advocated replacing it with an ethic that emphasizes strength and dominance.

This is the “Metaphysical Argument”: *Ethical opinions cannot be objectively true or false because there is no moral reality that they may match or fail to match.* This is the deep way in which ethics differs from science. Science describes a reality that exists independent of observers. If conscious beings ceased to exist, the world would otherwise be unchanged—it would still be there, and it would still be just as science describes it. But if there were no conscious beings, there would be no moral dimension to reality at all.

We may summarize the argument like this:

- (1) There are objective truths in science because there is an objective reality—the physical world—which science describes.
- (2) But there is no moral reality comparable to the reality of the physical world. There is nothing “there” for ethics to describe.
- (3) Therefore, there are no objective truths in ethics.

Is this argument sound? It is true, I think, that there is no moral reality comparable to the reality of the physical world. However, it does not follow from this that ethics has no objective basis. After all, an inquiry might be objective in two ways:

1. It may be objective because there is an independent reality that it describes correctly or incorrectly. Science is objective in this sense.
2. An inquiry may be objective because there are reliable methods of reasoning that determine truth and falsity in its domain. Mathematics is objective in this sense. Mathematical results are objective because they are provable by the relevant kinds of arguments.

Ethics is objective in the second sense. We do not discover whether an ethical opinion is true by comparing it to some sort of “moral reality.” Instead, we discover what is right by examining the reasons, or arguments, that can be given on both sides of an issue—the right thing to do is whatever there are the best reasons for doing. For ethics to be objective, it is enough that we can identify and evaluate reasons for and against ethical judgments.

11.4. The Importance of Human Interests

The preceding discussion should have dispelled most of your doubts about the objectivity of ethics. But some nagging doubts might remain, for good reason. We have not yet gotten to the bottom of things.

Every inquiry, whether in science, mathematics, or ethics, involves reasoning: We gather data, marshal arguments, and draw conclusions. But reasoning cannot go on forever. If I tell you that A is true, and you want to know why, I may cite B as my reason. If you call B into question, I may justify B by appealing to C. And so on. But at some point we come to the end of the line. Every chain of reasoning must end somewhere. This means that every argument ultimately appeals to some consideration that is simply taken for granted.

Scientific reasoning terminates when we reach simple facts about the physical world. We know, for example, that the galaxies are moving apart. How do we know this? Because of facts about the light that reaches the earth—in particular, facts about the red shift in the spectrum. How do we know what the red shift means? Because of many past observations and experiments. This example is simplified, but when we reach the simplest observed facts, we have reached the bedrock on which everything else rests. Mathematical reasoning is somewhat different in that it does not appeal to facts about the physical world. Instead, it relies on axioms, which may be taken as self-evidently true or may simply be assumed for the purposes of the proof.

Where do ethical arguments terminate? To what do they ultimately appeal? Let us look more closely at one of our previous examples. Smith is a bad man because, among other things, he is a habitual liar. This is a good reason for judging him to be morally deficient, we said, and so this fact forms part of a “proof” that he behaves unethically. But why is it bad to lie?

Lying is bad for several reasons. First, it is harmful to people. If I lie to you, and you believe me, then things can go badly for you. Suppose you ask me when the concert starts, and I say “Ten,” even though I know it begins at 7:30. You arrive at 9:40 only to find that you’ve missed half the show. Multiply this example many times over, and you will see why honesty is important. Second, lying is a violation of trust. When you believe me without checking on what I say, you are trusting me. So, if I lie to you, I am causing you harm in a special way, by taking advantage of your trust. This is why people feel so personally affronted when they’re lied to. Finally, the rule against lying is a fundamental social rule, in the sense that no society could exist without it. If we cannot assume that people will speak truthfully, then communication cannot take place; and without communication, society cannot exist.

Thus, the judgment that lying is wrong is not arbitrary. It has good reasons behind it. Suppose, though, you pushed further and wanted to know why it matters if people are harmed, or trust is violated, or society collapses. We could say a little more. We could point out that people are *worse off* when they base their decisions on false information and their trust is violated. We could point out that everyone would be worse off if people couldn’t live together in cooperative societies. But you persist: Why does it matter if people are worse off? Here we come to the end of the line. Ethical reasoning terminates in considerations about people being better or worse off—or, perhaps, in considerations about *any* sensitive creature being better or worse off—just as scientific reasoning terminates in simple observations about the physical world.

Some people conclude from this that ethics is subjective—after all, ethics is ultimately about the well-being of individuals, which is a subjective matter. Personally, I don’t think this is a good use of the word “subjective,” but what’s important is not what words we use, but what conclusions we draw from them. If ethics is subjective in the sense of being about individuals, it does not follow that ethics is arbitrary. Nor does it follow that people are free to accept whatever ethical judgments they like, or that one person’s opinions will always be as good as another’s. Ethics remains a matter of following reason, and it will still be objectively true that some things are good for people, and other things are bad for people. Ethical judgments can still be correct or incorrect. In these ways, the objectivity of ethics is secure.

Why Should We Be Moral?

Immorality is the morality of those who are having a better time.

—H. L. MENCKEN,
A MENCKEN CHRESTOMATHY (1956)

12.1. The Ring of Gyges

An ancient legend tells the story of Gyges, a poor shepherd who found a magic ring in a fissure opened by an earthquake. Gyges discovered that when he twisted the ring on his finger, he would become invisible. With this great power, Gyges could go anywhere and do anything he wanted, without fear of detection. He used the ring to enrich himself, taking what he wanted and killing anyone who got in his way. Eventually, he invaded the royal palace, where he seduced the queen, murdered the king, and seized the throne.

Glaucon tells this story in Book I of Plato's *Republic*. The story is meant to illustrate how behaving immorally can sometimes be to one's advantage. If Gyges had remained virtuous, he would have remained poor. By breaking the moral rules, however, he became rich and powerful. Considering this, why should Gyges care about the moral rules? For that matter, why should any of us be moral if it doesn't serve our own needs? Why tell the truth, if lying is more convenient? Why give money to charity, when you can spend it on yourself? Morality places restrictions on us that we may not like. So why shouldn't we just forget about it? Glaucon adds that, in his opinion, all of us would behave like Gyges, if we thought we could get away with it.

In what follows, we will take up the question of why we should be moral. But first, we need to understand the question. It is not a request for a justification of moral behavior. If it were,

then the answer would be easy. We could easily say why Gyges shouldn't have robbed and murdered his way to the throne. Robbery is taking things that do not belong to you, and murder inflicts a terrible harm on its victims. Similarly, it is easy to explain why we should be truthful or why we should give money to help the needy. Lying harms people, and hungry people need food more than people like us need money.

Such reasons, however, only determine what is right, and that is not the issue. Glaucon's challenge arises after the moral reasoning has been done. We may grant that it is morally right to respect people's lives and property. We may concede that it is right to tell the truth and help people. Glaucon's question is, *Why should we care about doing what is right?*

To answer this question, we must show that living morally is in our own best interests, and that will not be easy. On the surface, it looks like ethics is a hindrance in promoting one's own happiness. Of course, it may be a good thing for you if *other people* live ethically, for then they will respect your rights and be helpful to you. But if *you* are bound by moral constraints, that is another matter.

Can it be shown that accepting moral constraints really is in one's self-interest? We might be able to meet Glaucon's challenge, and we might not. It all depends on what morality is.

12.2. Ethics and Religion

One familiar idea is that right living consists in obedience to God's commands. On this conception, God has set out the rules we must obey, and he will reward those who follow his rules and punish those who do not.

If this were true, we could answer Glaucon's challenge as follows: We should be moral because otherwise God will punish us. Even if, like Gyges, we had the power of invisibility, we would still be visible to God, and so, ultimately, we could not get away with being wicked. In one familiar scenario, the righteous will spend eternity in heaven while the sinners will go to hell. Thus, any benefit you might gain from wrongdoing will be only temporary. In the long run, virtue pays.

That response, however, may not fully answer Glaucon's challenge. On the familiar scenario, God does not reward individuals in proportion to their virtue. Rather, God sorts

everyone into just two groups: the heaven-bound and the hell-bound. In the end, both the saints and those who barely scrape by wind up in the same heavenly paradise. Thus, it might be to your advantage to do just enough to get into heaven, but no more. And so, you can ask: Why should I be moral, when a little immorality will benefit me on earth without spoiling my chances at heaven?

On these assumptions, Glaucon's challenge can be partly answered: Even if we can get away with some immorality, we shouldn't risk hellfire by robbing and murdering our way to the throne. But in fact, much more virtue than that might be called for. In examining our own lives, we can never be too sure as to what God might be thinking. You never know when one more selfish act might tip the scales. Thus, we should probably be as good as we can be, to increase our chances of eternal bliss.

Let's now discuss, more generally, the moral theory that puts God in the center of things. According to the Divine Command Theory, an action is morally right if, and only if, it is commanded by God. Is this theory correct? If we tried to live by it, we would encounter enormous practical difficulties. How can we know what God commands? Some people claim that God has told them what he wants us to do. But why should we trust those people? Hearing voices can be a sign of schizophrenia, and anyway, they might be lying. They might be saying they've spoken to God merely to get attention, or to feel important, or to persuade their followers to send them money.

Others rely, more modestly, on Scripture and Church tradition for guidance. But those sources are notoriously difficult to interpret. They give vague and sometimes contradictory instructions. So, when people consult these authorities, they typically rely on their own judgment to sort out what seems acceptable. For example, they may cite the passage in Leviticus that condemns homosexuality while ignoring the passage that requires you to wash your clothes if you touch anything a menstruating woman has sat upon.

But these are just practical difficulties. It may still seem plausible that God's commands provide the ultimate basis for ethics: God's saying that something is wrong is what *makes it* wrong. Many religious people think it would be sacrilege not to accept this view. Socrates, however, gave a powerful argument against it.

In the *Euthyphro*, Socrates considers whether “right” can be the same as “what the gods command.” Socrates accepted that the gods exist and that they may issue instructions. But he showed that this cannot be the ultimate basis of ethics. As he said, we need to distinguish two possibilities: Either the gods have good reasons for their instructions, or they do not. If they do not, then their commands would be arbitrary—the gods would be like petty tyrants who demand that we do this and that even though there is no good reason for it. This is an impious view that religious people will not want to accept. On the other hand, if the gods do have good reasons for their instructions, then there must be a standard of rightness independent of their commands—namely, the standard which they themselves adhere to in deciding what to require of us.

It follows, then, that the rightness or wrongness of actions cannot be understood merely in terms of their conformity to divine commands. 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.

The same is true of sacred texts. Nothing can be morally right or wrong simply because a book says so. If the book’s decrees are not arbitrary, then there must be some reason for them. For example, the Bible says we should not lie about our neighbors—we should not “bear false witness” against them. Is this rule arbitrary? Certainly not. Lying causes harm and violates the trust that others have in us, and lying about our neighbors is insulting to them. *That’s* why lying is wrong. The reason is not, “because the Bible says so.” Similarly, we may ask why homosexuality is condemned. Are there good reasons for this pronouncement? If so, then those reasons will give the real explanation of why homosexuality is wrong. If not, then homosexuality isn’t wrong, and the Bible’s condemnation is unjustified.

These problems make the Divine Command Theory implausible, but they do not refute the separate idea that, if God punishes wrongdoing, then we have good reason to be moral. This idea so impressed the great German philosopher Immanuel Kant (1724–1804) that he made it into an argument for God’s existence. Kant reasoned that if God does not exist, then the universe is morally incomplete, because virtue will go unrewarded and wickedness will go unpunished. This thought was

intolerable to him, so he concluded that God exists. Even great philosophers, it seems, can indulge in wishful thinking.

12.3. The Social Contract

In the seventeenth century, with the rise of modern science, philosophy became an increasingly secular enterprise. Since then, there has been a rough consensus that ethics must be understood as a human phenomenon—as the product of human needs, interests, and desires.

Thomas Hobbes (1588–1679) was the first great modern thinker to offer a secular basis for ethics. Hobbes assumed that “good” and “bad” are what we call the things that we like or dislike. Thus, when you and I like different things, we regard different things as good or bad. However, Hobbes said, we are all essentially alike. We are all self-interested creatures who want to live as well as possible. This is the key to understanding ethics. Ethics arises when people realize what they must do to live as well as they can.

Hobbes points out that each of us is enormously better off living in a cooperative society than trying to make it on our own. The benefits of social living go far beyond companionship. Social living makes possible schools, hospitals, and highways; houses with electricity and central heating; airplanes and cell phones; books, magazines, and websites; movies, opera, and football; and science, engineering, and agriculture. Without social cooperation, we would lose all of this. Therefore, it benefits each of us to establish and maintain a cooperative society.

But it turns out that a 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:

- If people couldn’t be relied on to tell the truth, then no one would care what anyone said. Communication would be impossible. And without communication, society would collapse.
- If people didn’t keep their promises, then there could be no division of labor—workers could not count on getting paid, retailers could not rely on their suppliers, and so on—and the economy would collapse. There could be no business, no building, no agriculture, and no medicine.

- Without assurances against assault, murder, and theft, no one could feel secure. Everyone would have to be constantly on guard, and no one could trust their neighbors.

Thus, to obtain the benefits of social living, we must strike a bargain with one another, whereby each of us agrees to obey these rules. This “social contract” is the basis of morality. Thus, morality may be understood as *the rules that a self-interested person would agree to obey, provided that others agree to obey them as well.*

Why the Social Contract Theory Is Attractive. This way of thinking about 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. Instead, it is about doing what it takes to make social living possible.

Second, the social contract approach gives us a sensible way of determining what our moral duties are. We are morally required to do the things that make social living possible. Otherwise, we may do what we like. Unfortunately, when many people hear the word “morals,” they think about an attempt to restrict their sex lives. But an ethic based on the social contract would have little interest in people’s personal affairs.

Third, the Social Contract Theory gives a plausible answer to Glaucon’s question: Why should we behave morally when it is not to our advantage? Well, it *is* to our advantage to live in a society in which people behave morally. Thus, it is rational for us to accept moral restrictions on our conduct, because the overall system benefits us. There is, however, a problem with this answer.

The Problem of the Free Rider. Schemes of social cooperation always face the “free rider” problem. A free rider is someone who benefits from a cooperative arrangement without contributing to it. Suppose the home owners in my neighborhood chip in to pay for a streetlight. I want the light as much as anyone else does, but I refuse to chip in, knowing that they will go ahead without me. The light is installed, and I get to use it for free. That makes me a free rider.

In any cooperative society, there will be some free riders—individuals who benefit from living in the society, but who don’t

follow the rules which make the society possible. Glaucon's challenge may therefore be reformulated: Why not be a free rider? Each of us already enjoys the benefits of living in a stable society. So, why shouldn't we break the rules if we think we can get away with it?

The free-rider problem for ethics can be solved, but only partially. The partial solution goes like this: Each of us has good reason not only to encourage others to obey the social rules but also to make it as hard as possible for them to break those rules. Take the rule against murder, for example. You don't merely want to encourage others not to murder you. You want a situation in which *no one can get away with* murdering you. Each of us has good reason to support the creation and maintenance of a social system in which other people could not murder us. To accomplish this, we establish laws and other methods of enforcement. But in doing so, we create a situation in which *we* cannot get away with murder, either.

I said that this is a "partial" solution, for two reasons. First, we use the power of the law to enforce the rules against murder, theft, and other grave offenses, but not all social rules are suitable for legal enforcement. Rules of ordinary decency must be enforced in "the court of public opinion" rather than in the court of law. The penalty for lying, for example, is only that people will get angry at you. Although everyone wants to avoid prison, not everyone minds upsetting the people around them. Thus, informal methods of enforcement are bound to be only partly effective.

Second, no mechanisms, formal or informal, are going to be perfect. It is easy to get away with the occasional lie. One might even get away with the occasional murder. Glaucon's challenge remains: Why should you obey the moral rules when you think you can get away with breaking them?

12.4. Morality and Benevolence

The Social Contract Theory does not assume that people are altruistic. Each person can be motivated to obey the social rules out of simple self-interest. However, people are not entirely selfish. Human beings have at least some benevolent feelings, if only for family and friends. We have evolved as social creatures just as surely as we have evolved as creatures with lungs. Caring for our kin and members of our local group is as natural for us as breathing.

If humans do have some degree of natural altruism, does this have any significance for ethics? David Hume (1711–1776), the great Scottish philosopher, thought so. Hume agreed with Hobbes that our moral opinions are expressions of our feelings, but he did not believe that our feelings are merely self-centered. He believed that we also have “social sentiments”—feelings that connect us with other people and make us care about them. That is why, Hume says, 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 . . . be decided with greater certainty than by ascertaining . . . 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 our moral duties: *We should always try to produce the greatest possible benefit for everyone who will be affected by our action.*

This Principle of Utility is deceptively simple. It actually combines three ideas. First, in determining what to do, we should be guided by the consequences of our actions—we should do whatever can be reasonably expected to have the best outcome. Second, in determining which consequences are best, we should care only about the benefits and harms that would be caused—we should do whatever will cause the greatest benefits and the least significant harms. And third, the Principle of Utility assumes that each person’s happiness is as important as anyone else’s.

Although Hume suggested this idea, two philosophers from England pursued it in greater detail. Jeremy Bentham (1748–1832) led a movement to reform the laws of Britain along utilitarian lines. The Benthamites were remarkably successful in advancing such causes as prison reform and restrictions on the use of child labor. John Stuart Mill (1806–1873), the son of a Benthamite, gave the theory its most popular and influential defense in his book *Utilitarianism*.

The utilitarian movement attracted critics from the outset. It was an easy target because it ignores conventional religious notions. Morality, according to the Utilitarians, has nothing to do with obedience to God or gaining credit in heaven. Rather, the point is just to make life in this world as comfortable and as happy as possible. Thus, 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 rejected many conventional moral notions. Bentham argued, for example, that the purpose of the criminal justice system cannot be understood in the traditional way as “paying back” offenders 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. Today these ideas are familiar, but only because the utilitarian movement was so successful. 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 nor sex nor social class makes a difference to one’s moral status. Mill himself wrote a book called *The Subjection of Women* that became a classic of the nineteenth-century suffragist movement.

Finally, Utilitarianism was controversial because it has no use for “absolute” moral rules. According to Utilitarianism, the traditional moral rules are merely “rules of thumb”—they’re good pieces of advice, but they admit of exceptions. Whenever breaking a rule will have better results than following the rule, we should break it. The rule against killing, for example, might be suspended if someone is dying of a painful illness and requests a painless death. Moreover, the Utilitarians regard some traditional rules as suspect, even as rules of thumb. Christian

moralists had traditionally said that masturbation is evil, but from a utilitarian point of view, masturbation is a good thing. A more serious matter is the religious condemnation of homosexuality, which has resulted in misery for countless people. Utilitarianism implies that if an activity makes people happy, without anyone being harmed, it cannot be wrong.

Utilitarianism says that our moral duty is to promote the general happiness. Why should we do that? Mill echoes Glaucon's challenge when he says, "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 saw only one possible reason for accepting this or any other moral standard. The "internal sanction" of morality, he said, must always be "a feeling in our minds." And the kind of ethic we accept, he thought, will depend on the nature of our feelings. If human beings have "social feelings," then Utilitarianism 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 has implications that conflict with traditional morality. Much the same could be said about the Social Contract Theory. In most of the practical matters that we have mentioned—criminal punishment, racial discrimination, women's rights, euthanasia, homosexuality—the two theories agree. But they differ dramatically on one issue. Utilitarians, but not social contract theorists, think that we have an extensive moral duty to help others.

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). Around 8.8 million children under the age of 5 die each year, mostly from preventable causes such as pneumonia, diarrhea, malaria, and birth asphyxia. By giving your money to UNICEF

and making do with your old carpet, you could provide life-saving medical care for dozens of children. From the point of view of utility—seeking the best overall outcome for everyone—there is no doubt that you should give the money to UNICEF. Obviously, the medicine will help the children far more than the new rug will help you.

The Social Contract Theory takes a different approach. If morality rests on an agreement between people—an agreement we enter into to promote our own interests—what would the agreement say about helping other people? It depends on whether those other people can help us. If they can, then we could benefit from an agreement to help each other. If not, then we have no reason to accept any restrictions on our conduct.

From this point of view, we would have no reason to accept a general duty to provide aid to children in foreign countries. Jan Narveson, a social contract theorist, says that we needn't "go very far out of our way to be very helpful to those we don't know and may not particularly care for." And, Narveson asks,

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 help strangers or not? Here is a thought experiment that might help. Suppose there are two buttons on my desk, and I must choose which to press. By pressing button A, I can provide my daughter with a memorable party; by pressing button 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 daughter? What would your "conscientious feelings" tell you? Mill believed that one's conscientious feelings—the feelings that prevail after everything has been thought through—determine one's obligations. He believed 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. Their argument goes like this: It may be true 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 use those, we will realize that there are no relevant differences between us, those close to us, and strangers. Strangers have needs and interests, just like we do. Thus, no one should take his or her own well-being to be especially important. Peter Singer (1946–), a utilitarian philosopher, writes:

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 centre 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. Morally, everyone counts equally, even strangers in foreign countries who have no ability to help us or harm us.

Conclusion. We have reached no firm conclusions about what morality requires of us. However, it looks like morality requires us to care about the needs and interests of total strangers. But why should we do that? Glaucon’s challenge remains: Why should we do what’s right, if there’s nothing in it for us?

A religious outlook that includes a belief in the afterlife could help us answer Glaucon’s question. Otherwise, we can appeal to the fact that everyone benefits from a social arrangement in which the moral rules are acknowledged and enforced. But the free-rider problem cannot be completely solved; people can always get away with a certain amount of bad behavior. In the end, we can only hope that people’s behavior will be guided by what Mill called “the conscientious feelings of mankind.”

This may seem like a feeble conclusion. Yet “the conscientious feelings of mankind” are a powerful force, made stronger by education and the advancement of civilization. Still, we must concede that, if people can get away with wrongdoing and genuinely don’t care about others, then nothing will stop them. We could remind them of all the reasons why their actions would be wrong and remind them that their interests don’t matter more than the interests of other people. But that will only prove that their actions *are* wrong. To stop them from behaving badly, something more is required, namely, that they care about those reasons.