



Philosophical Ethics

Before embarking on analysis of the ethical issues surrounding computer and information technology, it will be helpful to discuss the nature of ethical analysis, and to become familiar with some traditional ethical concepts and theories. This chapter shows how ethical analysis can proceed so as to produce insight and better understanding. The chapter also explains concepts and theories that philosophers have found particularly useful in discussing ethical issues.

We often overhear or participate in discussions of ethical issues. Think, for example, of the heated discussions you have heard about government restrictions on individual freedom (e.g., censorship of the Internet, the right to assisted suicide). Or think of discussions about abortion, affirmative action, and the distribution of wealth in our society. Often when individuals are asked to explain why they think a behavior or policy is wrong, they have difficulty articulating their reasons. Sometimes it seems that individuals who are expressing moral opinions are simply reacting as they think most people in their society react or they espouse ideas they heard friends or relatives espouse. Many who have fairly strong moral beliefs have only a very vague sense of why the behavior or policy is unfair or irresponsible or harmful. These unexamined beliefs can be the starting place for ethical analysis, though it is important to understand that they are only starting places.

Discussions at this level may quickly end unresolved because the individuals involved are not able to provide good reasons for believing as they do. It is difficult or impossible to discuss the issues rationally, let alone resolve them. If discussion stays merely at the level of statements of belief, discussants will walk away thinking that everyone is entitled to his or her own opinion and there is no point talking about ethics, except perhaps to see where others stand. Discussants won't have learned anything or come to understand the ethical issues any better.

This book is an undertaking in philosophical analysis, and philosophical analysis proceeds on the premise that we must examine the reasons we have for our moral or ethical beliefs. In philosophical, ethical analysis the reasons for moral beliefs are articulated, and then critically evaluated. The reasons you give for holding an ethical belief or taking a position on an ethical issue can be thought of as an argument for a claim. The argument has to be "put on

the table,” and once there, it can be evaluated in terms of its plausibility, coherence, and consistency. Once stated, we can ascertain whether the argument does, indeed, support the claim being made or the position being taken.

This critical evaluation is often done in the context of trying to convince someone to reject a position, or to adopt another position, but it may also be done simply to explore a claim. When you critically evaluate the argument supporting a claim, you come to understand the claim more fully. A critical examination of the underpinnings of moral beliefs sometimes leads to a change in belief, but it may also simply lead to stronger and better understood beliefs.

In philosophical analysis, not only must you give reasons for your claims, you are also expected to be consistent from one argument or topic to the next. For example, instead of having separate, isolated views on abortion and capital punishment, philosophical analysis would lead you to recognize that both your views on abortion and your views on capital punishment rest on a claim about the value of human life and what abrogates it. Philosophical analysis would lead you to inquire whether the claim you made about the value of human life in the context of a discussion of capital punishment is consistent with the claim you made about the value of human life in the context of a discussion of abortion. If the claims appeared to be inconsistent from the one context to the next, then you would be expected to change one of your claims or provide an account of how the two positions can be understood as consistent. In other words, you would show that seemingly inconsistent views are in fact consistent.

Philosophical analysis is an ongoing process. It involves a variety of activities. It involves expressing a claim and putting forward an argument or reasons for the claim, and it involves critical examination of the argument. If the argument does not hold up to critical examination, then it might be reformulated into a revised argument, perhaps rejecting aspects of the original argument but holding on to a core idea. The revised argument, then, has to be critically examined, and so on, with ongoing reformulation and critique. Philosophers often refer to this process as a *dialectic* (which is related to the word *dialogue*). We pursue an argument to see where it goes and to find out what you would have to know or assert to defend the argument and establish it on a firm footing.

In addition to moving from claims to reasons and arguments, and from one formulation of an argument to another, better formulation, the dialectic also moves back and forth from cases to principles or theory. To illustrate, take the issue of euthanasia. Suppose you start out by making the claim that euthanasia is wrong. You articulate a principle as the reason for this claim. Say, the principle is that human life has the highest value and, therefore, human life should never be intentionally ended. You might then test this principle by seeing how it applies in a variety of euthanasia cases. For example, is it wrong to use euthanasia when the person is conscious but in extreme pain? When the person is unconscious and severally brain damaged? When the person is terminally ill? When the person is young or elderly? Since your principle concerns the value of human life, it has implications beyond the issue of euthanasia. Hence, you might

also test it by applying it to completely different types of cases. Is the intentional taking of human life wrong when it is done in a war situation? Is intentional killing wrong when it comes to capital punishment? Given your position on these cases, you may want to qualify the principle or you may hold to the principle and change your mind about the cases. For example, after seeing how the principle applies in various cases, you may want to qualify it so that you now assert that one should never intentionally take a human life *except* in self-defense or *except* when taking a life will save another life. Or you might reformulate the principle so that it specifies that the value of human life has to do with its quality. When the quality of life is significantly and permanently diminished, while it is still not permissible to intentionally kill, it is morally permissible to let a person die.

The dialogue continues as the dialectic leads to a more and more precise specification of the principle and the argument. The process clarifies what is at issue and what the possible positions are. It moves from somewhat inchoate ideas to better and better arguments, and more defensible and better-articulated positions.

The dialectic (from an initial belief to an argument, from argument to better argument, and from theory to case, and back) does not always lead to definitive conclusions or unanimous agreement. Therefore, it is important to emphasize that understanding can be improved, progress can be made, even when one has not reached definitive conclusions. Through the dialectic we learn which arguments are weaker and stronger and why. We come to understand the ideas that underpin our moral beliefs. We develop deeper and more consistent beliefs and we come to understand how moral ideas are interrelated and interdependent.

As you will see in a moment, a familiarity with traditional ethical theories will help in articulating the reasons for many of your moral beliefs. Ethical theories provide frameworks in which arguments can be cast. Moreover, ethical theories provide some common ground for discussion. They establish a common vocabulary and frameworks within which, or against which, ideas can be articulated.

DISTINGUISHING DESCRIPTIVE AND NORMATIVE CLAIMS

In any discussion of ethics, it is important to recognize the distinction between descriptive and normative claims. In a sense and partly, this is the distinction between facts and values, but the matter of what counts as a fact is very contentious in philosophy. So, it will be better to stay with the terms descriptive and normative. *Descriptive* statements are statements that describe a state of affairs in the world. For example, “The car is in the driveway.” And “Georgia is south of Tennessee.” In addressing ethical issues and especially the ethical issues surrounding computer and information technology, it is quite common to hear seemingly factual statements about human beings. The following are descriptive statements: “Such and such percentage of the people surveyed admitted to

having made at least one illegal copy of computer software.” “The majority of individuals who access pornographic Web sites are males between the ages of 14 and 35.” “Such and such percentage of U.S. citizens use the Internet to obtain information on political candidates.” “In all human societies, there are some areas of life that are considered private.” These statements describe what human beings think and do. They are *empirical* claims in the sense that they are statements that can be verified or proven false by examining the state of affairs described. To be sure, it may not be easy to verify or disconfirm claims like these, but in principle it is possible. Observations can be made, surveys can be administered, people can be asked, and so on.

Social scientists gather empirical data and report their findings, both on moral and nonmoral matters. When it comes to morality, psychologists and sociologists might do such things as identify the processes by which children develop moral concepts and sensibilities. Or they may measure how individuals value and prioritize various goods such as friendship, privacy, and autonomy. When anthropologists go to other cultures, they may describe complex moral rules in that culture. They are describing lived and observed moral systems. Similarly, historians may trace the development of a particular moral notion in an historical period.

All of these social scientific studies are descriptive studies of morality; they examine morality as an empirical phenomenon. They don't, however, tell us what is right and wrong. They don't tell us what people *should* think or do, only what people, in fact, think and do.

In contrast, philosophical ethics is *normative*. The task of philosophical ethics is to explore what human beings ought to do, or more accurately, to evaluate the arguments, reasons, and theories that are proffered to justify accounts of morality. Ethical theories are prescriptive. They try to provide an account of why certain types of behavior are good or bad, right or wrong. Descriptive statements may come into play in the dialectic about philosophical ethics, but normative issues cannot be resolved just by pointing to the facts about what people do or say or believe. For example, the fact (if it were true) that many individuals viewed copying proprietary software as morally acceptable would not make it so. The fact that individuals hold such a belief is not an argument for the claim that it is morally permissible to copy proprietary software. You might wish to explore why individuals believe this to see if they have good reasons for the belief. Or you might wish to find out what experiences have led individuals to draw this conclusion. Still, in the end, empirical facts are not alone sufficient to justify normative claims. Figuring out what is right and wrong, what is good and what is bad, involves more than a descriptive account of states of affairs in the world.

The aim of this book is *not* to describe how people behave when they use computers. For this, the reader should consult social scientists—sociologists, anthropologists, political scientists, and psychologists. Rather the aim of this book is to help you understand how people ought to behave when they use computers and what rules or policies ought to be adopted with regard to computer and information technology.

ETHICAL RELATIVISM

We can begin our examination of ethical concepts and theories by examining a prevalent, often unexamined moral belief. Many believe that “ethics is relative.” This seems like a good starting place. This claim can be examined carefully and critically. We can begin by formulating the idea as a theory consisting of a set of claims backed by reasons.

The idea of ethical relativism seems to be something like this: “What is right for you may not be right for me,” or “I can decide what is right for me, but you have to decide for yourself.” When we take this idea and formulate it into a more systematic account, it seems to encompass a negative claim (something that it denies), and a positive claim, (something it asserts). The negative claim appears to be: “There are no universal moral norms.” According to this claim, there isn’t a single standard for all human beings. One person may decide that it is right for him to tell a lie in certain circumstances, another person may decide that it is wrong for her to tell a lie in exactly the same circumstances, and both people could be right. So, the claim that “right and wrong are relative” means in part that there are no universal rights and wrongs.

The positive claim of ethical relativism is more difficult to formulate. Sometimes ethical relativists seem to be asserting that right and wrong are relative to the individual, and sometimes they seem to assert that right and wrong are relative to the society in which one lives. I am going to focus on the latter version, and on this version the relativist claims that what is morally right for me, an American living in the twenty-first century, could be different than what is right for a person living, say, in Asia in the fifth century. The positive claim of relativism is that right and wrong are relative to your society.

Ethical relativists often cite a number of descriptive facts to support these claims:

1. They point to the fact that cultures vary a good deal in what they consider to be right and wrong. For example, in some societies, infanticide is acceptable while in other societies it is considered wrong. In some societies, it is considered wrong for women to go out in public without their faces being covered. Polygamy is permissible in some cultures; in others it is not. Examples of this kind abound.
2. Relativists also point to the fact that the moral norms of a given society change over time so that what was considered wrong at one time, in a given society, may be considered right at another time. Slavery in America is a good example of this since slavery was considered morally permissible by many in the United States at one time, but is now illegal and almost universally considered impermissible.
3. Relativists also point to what we know about how people develop their moral ideas. We are taught the difference between right and wrong as children, and what we come to believe is right or wrong is the result of our upbringing. It depends on when, where, how, and by whom we were raised. If I had been born in certain Middle Eastern countries, I might believe that it is wrong for a woman to appear in public without her face covered. Yet because I was raised in the

United States in the twentieth century, by parents who had Western ideas about gender roles and public behavior, I do not believe this. Of course, parents are not the only determinant of morality. A person develops moral ideas from the experiences he or she has in school, at work, with peers, and so on.

It is useful to note that we have already made progress simply by clearly and systematically formulating the idea of ethical relativism, an idea you may have entertained or heard expressed, but never had a chance to examine carefully. Moreover, we have been able to identify and articulate some reasons thought to support ethical relativism. With the idea and supporting evidence now “on the table,” we can carefully and critically examine them.

The facts which ethical relativists point to cannot be denied. For example, I would not want to take issue with the claims that:

1. There is and always has been a good deal of diversity of belief about right and wrong.
2. Moral beliefs change over time within a given society.
3. Social environment plays an important role in shaping the moral ideas you have.

However, there does seem to be a problem with the connection between these facts and the claims of ethical relativism. Do these facts show that there are no universal moral rights or wrongs? Do they show that right and wrong are relative to your society?

On more careful examination, it appears that the facts cited by ethical relativists do not support their claims. To put this another way, we can, without contradiction, accept the facts and still deny ethical relativism. The facts do not necessitate that there are no universal moral standards or that ethics is relative. Lest there be no confusion, you should recognize that “ethics is relative” could be interpreted either as an empirical or a normative claim. As an empirical claim, it asserts that ethical beliefs vary; as a normative claim it asserts that right and wrong (not just beliefs about, but what is actually right and wrong) vary.

If we understand the claim “ethics is relative” to be a description of human behavior, then it *does* follow from the facts cited. Indeed, it is redundant of the facts cited, for as a description of human behavior, it merely repeats what the facts have said. Ethical beliefs vary. Individuals believe different things are right and wrong depending on how and by whom they have been raised and where and when they live.

On the other hand, if we understand “ethics is relative” to be a normative claim, a claim asserting the negative and/or positive parts of ethical relativism, then it is not redundant, and the facts do not support the claims. Here the leap from facts to conclusion is problematic for a number of reasons. For one, the argument goes from a set of “is” claims to an “ought” claim and the ought-claim just doesn’t follow (in a straightforward way) from the is-claims.

The argument goes like this: “People do a; people do b; people do c; and therefore people ought to do x.”

Moreover, the facts are compatible with the opposite conclusion. That is, it is possible that a universal moral code applies to everyone even though some or all fail to recognize it. Centuries ago when some people believed the earth was flat and others claimed that it was round, the earth’s shape was not relative. The fact that there is diversity of opinion on right and wrong does not tell us anything about whether right and wrong are relative. The facts are compatible both with the claim that there is no universal right and wrong and with the claim that there is a universal right and wrong.

Taking this one step further, let’s consider the fact that our moral beliefs are shaped by our social environment. While it is true that our moral beliefs are shaped by our social environment, this says nothing about the rightness or wrongness of what we believe. Racism and sexism are good examples of moral attitudes we may acquire from our environment but which turn out on reflection to be unjustifiable (bad) ideas.

We must also be careful about what is inferred from the fact that there is diversity in moral beliefs. This diversity may be misleading; that is, it may be superficial rather than deep. Relativists seem to be focusing on specific practices and there is still the possibility that universal norms underlie these. Moral principles such as “never intentionally harm another person” or “always respect human beings as ends in themselves” are of such generality that they could be operative in many or all cultures but expressed in different ways. What is meant by “harm,” “respect,” and “human being” may vary although there is some principle to which all people adhere. So, it is possible that there are some universal principles at work, but they are hidden from sight due to the diversity of expression or interpretation of the principle.

Social scientists have certainly tried to find patterns within the apparent diversity. Some have asserted, for example, that all cultures have prohibitions on incest or, more recently, that while there is a great deal of diversity about what is considered private, all cultures consider some aspect of the lives of individuals private.

Even so, while such patterns have important implications for the study of ethics, we have to remember that establishing patterns across cultures is descriptive, and it is another matter to determine what these claims imply about how people ought to behave. In a moment, when we examine utilitarianism, we will see an example of a very general normative principle that is compatible with a diversity of practices. Utilitarianism is a form of consequentialism, and such theories assert that individuals should always do what will maximize good consequences. Individuals in quite different situations may be doing very different things but all in accordance with this same principle.

In any case, the facts pointed to by relativists do not support their claim that there are no universal moral rights and wrongs. Nor do the facts cited support the ethical relativist’s claim that right and wrong are relative to

one's society. Pointing to what people believe to be right and wrong tells us nothing about what is right or wrong. The fact that people behave in accordance with the norms of their society is not evidence for the claim that they ought to.

It is important to keep in mind that the criticism I have just made of the ethical relativist's argument does not establish that there are universal rights and wrongs. The criticisms show only that the arguments ethical relativists might put forward to support their position do not work. You may be able to come up with a different argument on behalf of ethical relativism, and then your argument would have to be carefully and critically examined.

Before you try to defend ethical relativism, however, there are some serious problems with the theory and you ought to be aware of these. Ethical relativism, as I have formulated it, appears to be self-contradictory. The negative and positive claims appear to contradict each other. In saying that right and wrong are relative to one's society, ethical relativists seem to be asserting that one is bound by the rules of their society. The relativist seems to be saying that what is right for me is defined by my society, and what is right for a member of an African tribe is what is set by the standards of her or his tribe. It would seem, then, that I ought to do what is considered right in my society, and everyone else ought to do what is considered right in their society. Notice, however, that if this is what ethical relativists mean, they are affirming a universal moral principle. On the one hand, they deny that there are universal rights and wrongs, and, on the other hand, they assert one. If I have accurately depicted ethical relativism, then it appears to be an utterly incoherent (self-contradictory) theory.

If this were a book about ethical relativism alone, I would try to resurrect the theory by reformulating its claims and bringing in other arguments to support it. All I will do instead is to point to what I think is an important moral motive buried in relativism. Often what ethical relativists are trying to do is make the point that no one should denigrate, ridicule, and disrespect people who have beliefs that are different from their own. In other words, you shouldn't judge people from other times or places by the standards of your own morality. It is arrogant, relativists might say, to believe that you as an individual or a member of a particular society have the correct moral views and that anyone who doesn't agree with you is wrong. Such relativists would argue that we ought to respect people with moral beliefs different from our own.

This seems an important and worthy point that some relativists want to make. Still, it should be noted that to take this position is, again, to take a universal position. You are claiming that "everyone ought" to adopt a position which might be characterized as tolerance or respect for others.

So, it would seem that we cannot assert both that everyone ought to respect the views of others and at the same time hold that ethics is relative. If tolerance is the motive behind relativism, this motive has an implicit universal

character and that conflicts with relativism's claim that there is no universal right and wrong. To see the contradiction, consider the case of someone who lives in a society that does not believe in toleration. According to relativism, this person need not be tolerant of others. Relativism says right and wrong is relative to your society and in this person's society there is nothing wrong with being intolerant. Thus, it would seem that if underlying one's belief in relativism is the belief that everyone should be tolerant of the beliefs of others, relativism is not going to be an acceptable theory, at least not if it is formulated as I have formulated it.

Case Illustration

To see these and other problems with ethical relativism, consider a hypothetical case. Suppose, by a distortion of history, that computers were developed to their present sophistication in the late 1930s and early 1940s. World War II is in progress. You are a German citizen working for a large computer company. You are in charge of the sales division and you personally handle all large orders. You are contacted by representatives of the German government. The German government has not yet fully automated its operations (computers are still relatively new) and it wants now to purchase several large computers and several hundred smaller computers to be networked.

You read the newspapers and know how the war is proceeding so you have a pretty good idea of how the German government will use the computers. It is quite likely they will use the computers to help keep track of their troops and equipment, to identify Jews and monitor their activities, to build more efficient gas chambers, and so on. The question is, if you were an ethical relativist would it be permissible for you to sell the computers to Hitler and his government?

The question reveals some practical problems with relativism. Relativism specifies that what is right for you is what is considered right in your society. But, how do you figure out what the standards of your society are? Are the standards of your society what the political leaders say and do or what the majority in the society believe? If these are different, what should you do? To put this in another way, is Hitler necessarily abiding by the standards of his society or is he going against these? If he is going against these standards, then perhaps he is doing wrong and you would be doing wrong to support him. It may not be easy to tell whether Hitler is adhering to or rejecting the standards of his society. Hence, it may not be so easy to use relativism to guide your actions.

This leads to another problem with relativism. Suppose Hitler and most German citizens agree that Hitler's agenda is right. Nevertheless, you disagree. Relativism seems to rule out the possibility of resistance or rebellion in such a situation. If someone rebels against the standards of her society, it

would seem she is doing wrong for she is acting against relativism's claim that what is right for you is what is considered right in your society. Many of our greatest heroes, Socrates, Martin Luther King, Ghandi, even Jesus, would, on this account, be considered wrong or bad. They acted against the standards of their societies.

So, if Hitler and most Germans agreed that the German agenda was right, it would seem that you, as a relativist, would have to conclude that it is right for you to sell the computers to the German government (even if you personally objected to Hitler's agenda).

Now suppose that one of your friends from the United States or somewhere else finds out about the sale and asks you why you did this. What do you say? You answer: It was the right thing to do because it was consistent with the standards and beliefs in my society. From your friends perspective, this may seem a very feeble answer. The fact that some type of behavior is the standard in your society seems an inadequate moral reason for adopting the standard as your own. It doesn't seem a very good reason for acting in a certain way, especially when the act has significant negative consequences.

Summarizing what has been said so far about the problems with relativism, it suffers from three types of problems. First, the evidence that is used to support it, does not support it. Second, proponents cannot assert both the negative and the positive claims of relativism without inconsistency. By claiming that everyone is bound by the rules of his or her society, the ethical relativist makes a universal claim and yet the relativist claims there are no universal rights and wrongs. And, third, the theory, as the Hitler case illustrates, does not seem to help in making moral decisions. Relativism, at least as I have formulated it, does not help us figure out what to do in tough situations. It recommends that we adhere to the standards in our society and yet it doesn't help us figure out what these standards are. Moreover, doing something because it is the standard in your society does not seem a good reason for doing something.

Where do we stand now? It is important to note that we have made progress even though we have not formulated a moral theory that is defensible. Partly our progress is negative. That is, we have identified some arguments that don't work. At the same time, we have learned about some of the difficulties in taking a relativist position and are therefore in a better position to reformulate the theory. Perhaps, most important of all, we have seen the challenge of developing and defending ethical claims.

Our exploration of ethical relativism has hardly scratched the surface. You may want to reformulate ethical relativism so as to avoid some of the arguments given against it. You may want, for the time being, to take what might be called "an agnostic position." As an agnostic, you claim that you don't yet know whether there are universal rights and wrongs but you would also claim that you do not have sufficient reasons for ruling out the possibility either. You will

wait and see, keeping an open mind, and being on the alert for implausible and inconsistent claims.

UTILITARIANISM

Utilitarianism is an ethical theory claiming that what makes behavior right or wrong depends wholly on the consequences. In putting the emphasis on consequences, utilitarianism affirms that what is important about human behavior is the outcome or results of the behavior and not the intention a person has when he or she acts. On one version of utilitarianism, what is all important is happiness-producing consequences (Becker and Becker, 1992). Crudely put, actions are good when they produce happiness and bad when they produce the opposite, unhappiness. The term *utilitarianism* derives from the word *utility*. According to utilitarianism actions, rules, or policies are good because of their usefulness (their utility) in bringing about happiness.

Lest there be no confusion, philosophers are not always consistent in the way they use the terms utilitarianism and consequentialism. Sometimes, consequentialism is seen as the broadest term referring to ethical theories that claim that what makes an action right or wrong is the consequences and not the internal character of action. Utilitarianism is, then, a particular version of this type of theory with the emphasis specifically on happiness-producing consequences. That is the way I shall use these terms, though I warn readers that the distinction sometimes is made in just the opposite way, that is, with utilitarianism seen as the broadest theory and consequentialism as a particular form of utilitarianism.

In any case, in the version on which I will focus, the claim is that in order to determine what they should do, individuals should follow a basic principle. The basic principle is this: *Everyone ought to act so as to bring about the greatest amount of happiness for the greatest number of people.*

But, what, you may ask, is the “proof” of this theory? Why should each of us act to bring about the greatest amount of happiness? Why shouldn’t we each seek our own interest?

Intrinsic and Instrumental Value

Utilitarians begin by focusing on values and asking what is so important, so valuable to human beings, that we could use it to ground an ethical theory. They note that among all the things in the world that are valued, we can distinguish things that are valued because they lead to something else from things that are valued for their own sake. The former are called *instrumental* goods and the latter *intrinsic* goods. Money is a classic example of something that is instrumentally good. It is not valuable for its own sake, but rather has value as a means for acquiring other things. On the other hand, intrinsic goods

are not valued because they are a means to something else. They have qualities or characteristics that are valuable in themselves. Knowledge is sometimes said to be intrinsically valuable. So, is art because of its beauty. You might also think about environmental debates in which the value of nature or animal or plant species or ecosystems are said to be valuable independent of their value to human beings. The claim is that these things have value independent of their utility to human beings.

Having drawn this distinction between instrumental and intrinsic goods, utilitarians ask what is so valuable that it could ground a theory of right and wrong? It has to be something intrinsically valuable, for something which is instrumentally valuable is dependent for its goodness on whether it leads to another good. If you want x because it is a means to y , then y is what is truly valuable and x has only secondary or derivative value. Utilitarianism, as I am using the term, claims that happiness is the ultimate intrinsic good, because it is valuable for its own sake. Happiness cannot be understood as simply a means to something else. Indeed, some utilitarians claim that everything else is desired as a means to happiness and that, as a result, everything else has only secondary or derivative (instrumental) value.

To see this, take any activity that people engage in and ask why they do it. Each time you will find that the sequence of questions ends with happiness. Take, for example, your career choice. Suppose that you have chosen to study computer science so as to become a computer professional. Why do you want to be a computer professional? Perhaps you believe that you have a talent for computing, and you believe you will be able to get a well-paying job in computer science—one in which you can be creative and somewhat autonomous. Then we must ask, why are these things important to you? That is, why is it important to you to have a career doing something for which you have a talent? Why do you care about being well paid? Why do you desire a job in which you can be creative and autonomous? Suppose that you reply by saying that being well paid is important to you because you want security or because you like to buy things or because there are people who are financially dependent on you. In turn, we can ask about each of these. Why is it important to be secure? Why do you want security or material possessions? Why do you want to support your dependents? The questions will continue until you point to something that is valuable in itself and not for the sake of something else. It seems that the questions can only stop when you say you want whatever it is because you believe it will make you happy. The questioning stops here because it doesn't seem to make sense to ask why someone wants to be happy.

A discussion of this kind could go off in the direction of questioning whether your belief is right. Will a career as a computer professional make you happy? Will it really bring security? Will security or material possessions, in fact, make you happy? Such discussions center on whether or not you have chosen the correct means to your happiness. However, the point that utilitarians

want to make is that any discussion of what you should seek in life, and what is valuable, will not stop until we get to happiness.

It makes no sense, utilitarians argue, to ask why people value happiness. Happiness is the ultimate good. All our actions are directly or indirectly aimed at happiness. It is happiness for which we all strive. Utilitarians seem to believe that this is simply part of our human nature. Human beings are creatures who seek happiness. And, since happiness is the ultimate good, utilitarians believe that morality must be based on creating as much of this good as possible. Thus, all actions should be evaluated in terms of their “utility” for bringing about happiness.

According to utilitarianism, when an individual is faced with a decision about what to do, the individual should consider his or her alternatives, predict the consequences of each alternative, and choose that action which brings about the most good consequences, that is, the most happiness. So, the utilitarian principle provides a decision procedure. When you have to decide what to do, consider the happiness-unhappiness consequences that will result from your various alternatives. The alternative that produces the most overall net happiness (good minus bad) is the right action. To be sure, the right action may be one that brings about some unhappiness, but that is justified if the action also brings about so much happiness that the unhappiness is outweighed, or as long as the action has the least net unhappiness of all the alternatives.

Be careful not to confuse utilitarianism with *egoism*. Egoism is a theory that specifies that one should act so as to bring about the greatest number of good consequences for yourself. What is good is what makes “me” happy or gets me what I want. Utilitarianism does not say that you should maximize your own good. Rather, total happiness is what is at issue. Thus, when you evaluate your alternatives, you have to ask about their effects on the happiness of everyone. This includes effects on you, but your happiness counts the same as the happiness of others. It may turn out to be right for you to do something that will diminish your own happiness because it will bring about a marked increase in overall happiness.

The decision-making process proposed in utilitarianism seems to be at the heart of a good deal of social decision making. That is, legislators and public policy makers seem to seek policies that will produce good consequences, and they often opt for policies that may have some negative consequences but on balance, bring about more good than harm. Cost-benefit or risk-benefit analysis aims at quantifying net good consequences. This involves weighing the potential benefits of a project, such as construction of a new waste disposal plant, against the risks of harm in undertaking the project. It involves calculating and weighing the negative and positive effects of a project in deciding whether to go forward with it. In the case of a waste disposal plant, for example, we look at alternative ways to handle the waste, the various costs and benefits of each alternative, the good and bad effects of locating

the plant here or there, and so on. We balance the benefits of the plant against the risk of harm and other negative consequences to all those who will be affected.

Acts versus Rules

As mentioned earlier, there are several formulations of utilitarianism and proponents of various versions disagree on important details. One important and controversial issue of interpretation has to do with whether the focus should be on *rules* of behavior or individual *acts*. Utilitarians have recognized that it would be counter to overall happiness if each one of us had to calculate at every moment what all the consequences of every one of our actions would be. Not only is this impractical, because it is time consuming and because sometimes we must act quickly, but often the consequences are impossible to foresee. Thus, there is a need for general rules to guide our actions in ordinary situations.

Rule-utilitarians argue that we ought to adopt rules that, if followed by everyone, would, in the long run, maximize happiness. Take, for example, telling the truth. If individuals regularly told lies, it would be very disruptive. You would never know when to believe what you were told. In the long run, a rule obligating people to tell the truth has enormous beneficial consequences. Thus, “tell the truth” becomes a utilitarian moral rule. “Keep your promises,” and “Don’t reward behavior that causes pain to others,” are also rules that can be justified on utilitarian grounds. According to rule-utilitarianism, if the rule can be justified in terms of the consequences that are brought about from people following it, then individuals ought to follow the rule.

Act-utilitarians put the emphasis on individual actions rather than rules. They believe that even though it may be difficult for us to anticipate the consequences of our actions, that is what we should be trying to do. Take, for example, a case where lying may bring about more happiness than telling the truth. Say you are told by a doctor that tentative test results indicate that your spouse *may* be terminally ill. You know your spouse well enough to know that this knowledge, at this time, will cause your spouse enormous stress. He or she is already under a good deal of stress because of pressures at work and because someone else in the family is very ill. To tell your spouse the truth about the test results will cause more stress and anxiety, and this stress and anxiety may turn out to be unnecessary if further tests prove that the spouse is not terminally ill. Your spouse asks you what you and the doctor talked about. Should you lie or tell the truth? An act-utilitarian might say that the right thing to do in such a situation is to lie, for little good would come from telling the truth and a good deal of suffering (perhaps unnecessary suffering) will be avoided from lying. A rule-utilitarian would agree that good might result from lying in this one case, but in the long run, if we cannot count on people telling the truth (especially our spouses), more bad than good will come. Think of the anxiety that might arise if spouses routinely

lied to one another. Thus, according to rule-utilitarians, we must uphold the rule against lying; it would be wrong to lie.

Act-utilitarianism treats rules simply as “rules of thumb,” general guidelines to be abandoned in situations where it is clear that more happiness will result from breaking them. Rule-utilitarians, on the other hand, take rules to be strict. They justify moral rules in terms of the happiness consequences that result from people following them. If a rule is justified, then an act that violates the rule is wrong.

In either case, it should be clear that the utilitarian principle can be used to formulate a decision procedure for figuring out what you should do in a situation. In fact, many utilitarians propose that the utilitarian principle be used to determine the laws of a society. Laws against stealing, killing, breaking contracts, fraud, and so on can be justified on utilitarian grounds. Utilitarianism is also often used as a principle for evaluating the laws that we have. If a law is not producing good consequences or is producing a mixture of good and bad effects, and we know of another approach that will produce better net effects, then that information provides the grounds for changing the law. Punishment is a good example of a social practice that can be evaluated in terms of its utility. According to utilitarianism, since punishment involves the imposition of pain, if it does not produce some good consequences, then it is not justified. Typically utilitarians focus on the deterrent effect of punishment as the good consequence counterbalancing the pain involved.

Earlier I mentioned that utilitarianism might be said to capture part of the idea in relativism. According to utilitarianism, the morally right thing to do in a given situation will depend entirely on the situation. In one situation, it may be right to lie, in another situation in which the circumstances are different, it may be wrong to lie. Even rule-utilitarians must admit that the rule that will produce the most happiness will vary from situation to situation. A simple example would be to suppose a natural environment in which water is scarce. In such a situation, a rule prohibiting individuals from putting water in swimming pools and watering lawns would be justified. The rule would be justified because the alternative would lead to bad consequences. On the other hand, in a natural environment in which water is abundant, such a rule would not be justified.

So, even though utilitarians assert a universal principle, the universal principle has varying implications depending on the situation. This means that utilitarianism is consistent with varying laws and practices at different times or in different places depending on the specific circumstances.

Now that the fundamentals of utilitarianism have been explained, it is worth remembering, once again, that we are engaged in a dialectic. We have developed the idea of utilitarianism; we have made the case for the theory. The theory has been “put on the table,” so to speak. Even though it has been developed only in its most rudimentary form, the theory now needs to be critically scrutinized.

Critique of Utilitarianism

One of the important criticisms of utilitarianism is that when it is applied to certain cases, it seems to go against some of our most strongly held moral intuitions. In particular, it seems to justify imposing enormous burdens on some individuals for the sake of others. According to utilitarianism, every person is to be counted equally. No one person's unhappiness or happiness is more important than another's. However, since utilitarians are concerned with the total amount of happiness, we can imagine situations where great overall happiness might result from sacrificing the happiness of a few. Suppose, for example, that having a small number of slaves would create great happiness for a large number of individuals. The individuals who were made slaves would be unhappy, but this would be counterbalanced by significant increases in the happiness of many others. This seems justifiable (if not obligatory) according to utilitarianism. Another more contemporary example would have us imagine a situation in which by killing one person and using all their organs for transplantation, we would be able to save ten lives. Killing one to save ten would seem to maximize good consequences. Critics of utilitarianism argue that since utilitarianism justifies such practices as slavery and killing of the innocent, it has to be wrong. It is, therefore, unacceptable as an account of morality.

In defending the theory from this criticism, some utilitarians argue that utilitarianism does not justify such unsavory practices. Critics, they argue, are forgetting the difference between short-term and long-term consequences. Utilitarianism is concerned with all the consequences and when long-term consequences are taken into account, it becomes clear that such practices as slavery and killing innocent people to use their organs could never be justified. In the long run, such practices have the effect of creating so much fear in people that net happiness is diminished rather than increased. Imagine the fear and anxiety that would prevail in a society in which anyone might at any time be taken as a slave. Or imagine the reluctance of anyone to go to a hospital if there was even a remote possibility that they might be killed if by chance they were there when multiple organs were needed to save lives. The good effects of such practices could never counterbalance these bad effects.

Other utilitarians boldly concede that there are going to be some circumstances in which what seem to be repugnant practices should be accepted because they bring about consequences having a greater net good than would be brought about by other practices, that is, because they are consistent with the principle of utility. So, for example, according to these utilitarians, if there are ever circumstances in which slavery would produce more good than ill, then slavery would be morally acceptable. These utilitarians acknowledge that there may be circumstances in which some people should be sacrificed for the sake of total happiness.

In our dialogue about ethics, it is important to pick up on our strongly held moral intuitions for they are often connected to a moral principle or

theory. In the case of utilitarianism, the intuition that slavery is always wrong (or that it is wrong to kill the innocent for the sake of some greater good) points to an alternative moral theory. A concrete case will help us further understand utilitarianism and introduce a different theory, one that captures the moral intuition about the wrongness of slavery and killing the innocent.

Case Illustration

Not long ago, when medical researchers had just succeeded in developing the kidney dialysis machine, a few hospitals acquired a limited number of these expensive machines. Hospitals soon found that the number of patients needing treatment on the machines far exceeded the number of machines they had available or could afford. Decisions had to be made as to who would get access to the machines, and these were often life-death decisions. In response, some hospitals set up internal review boards composed of medical staff and community representatives. These boards were charged with the task of deciding which patients should get access to the dialysis machines. The medical condition of each patient was taken into account, but the decisions were additionally made on the basis of the personal and social characteristics of each patient: age, job, number of dependents, social usefulness of job, whether the person had a criminal record, and so on. The review committees appeared to be using utilitarian criteria. The resource—kidney dialysis machines—was scarce and they wanted to maximize the benefit (the good consequences) of the use of the machines. Thus, those who were most likely to benefit and to contribute to society in the future would get access. Individuals were given a high ranking for access to the machines if they were doctors (with the potential to save other lives), if they had dependents, if they were young, and so on. Those who were given lower priority or no priority for access to the machines were those who were so ill that they were likely to die even with treatment, those who were older, those who were criminals, those without dependents, and so on.

As the activities of the hospital review boards became known to the public, they were criticized. Critics argued that your value as a person cannot be measured by your value to the community. The review boards were valuing individuals on the basis of their social value and this seemed dangerous. Everyone, it was argued, has value in and of themselves.

The critique of this method for deciding who should live and who should die suggested a principle that is antithetical to utilitarianism. It suggested that each and every person, no matter what their social role or lot in life, has value and should be respected. To treat individuals as if they are a means to some social end seems the utmost in disrespect. And, that is exactly what a policy of allocating scarce resources according to social value does. It says, in effect, that people have value only as means to the betterment of society, and by that criteria some individuals are much more valuable than others.

The critics of distribution of kidney dialysis on the basis of social utility proposed as an alternative that scarce medical resources should be distributed by a lottery. In a lottery, everyone has an equal chance. Everyone counts the same. This, they argued, was the only fair method of distribution.

The kidney dialysis issue is just a microcosm of all medical resources. Doctors, medical equipment, and medical research are expensive and we have a finite amount of money to spend. Hence, lines have to be drawn—on what level of care goes to who, at what stage in their life, and so on. Distributive decisions have to be made.

The important point for our purposes is that the formulation of utilitarianism we have been considering leads to methods of distribution that seem to be unfair or unjust. So while the core idea in utilitarianism seems plausible (i.e., that everyone's happiness or well-being should be counted), utilitarianism does not seem to adequately handle the distribution of benefits and burdens. The criticism of the hospital review boards for distributing access to kidney machines according to social value goes to the heart of this criticism. Critics argue that people are valuable in themselves, not for their contribution to society. They argue that utilitarian programs are often unfair because in maximizing overall good, they impose an unfair burden on some individuals, and as such treat those individuals merely as means to social good.

I will now turn to an ethical theory that articulates the reasoning underlying the critique of utilitarianism. Before doing so, however, it is important to note that the dialectic could go off in a different direction. The debate about utilitarianism is rich and there are many moves that could be made in reformulating the theory and defending it against its critics. It is also important to note that whatever its weaknesses, utilitarianism goes a long way in providing a systematic account of many of our moral notions.

DEONTOLOGICAL THEORIES

In utilitarianism, what makes an action or a rule right or wrong is outside the action; it is the consequences of the action or rule that make it right or wrong. By contrast, deontological theories put the emphasis on the internal character of the act itself.¹ What makes an action right or wrong for deontologists is the principle inherent in the action. If an action is done from a sense of duty, if

¹ The term *deontology* is derived from the Greek words *deon* (duty) and *logos* (science). Etymologically, then, deontology means the science of duty. According to the *Encyclopedia of Philosophy*, its current usage is more specific, referring to an ethical theory which holds that "at least some acts are morally obligatory regardless of their consequences for human weal or woe." (Edwards, 1967)

the principle of the action can be universalized, then the action is right. For example, if I tell the truth (not just because it is convenient for me to do so, but) because I recognize that I must respect the other person, then I act from duty and my action is right. If I tell the truth because I fear getting caught or because I believe I will be rewarded for doing so, then my act is not morally worthy.

I am going to focus here on the theory of Immanuel Kant. If we go back for a moment to the allocation of dialysis machines, Kant's moral theory is applicable because it proposes what is called a categorical imperative specifying that we should never treat human beings merely as means to an end. We should always treat human beings as ends in themselves. Although Kant is not the only deontologist, I will continue to refer to him as I discuss deontology.

The difference between deontological theories and consequentialist theories was illustrated in the discussion of allocation of dialysis machines. Deontologists say that individuals are valuable in themselves, not because of their social value. Utilitarianism is criticized because it appears to tolerate sacrificing some people for the sake of others. In utilitarianism, right and wrong are dependent on the consequences and therefore vary with the circumstances. By contrast, deontological theories assert that there are some actions that are always wrong, no matter what the consequences. A good example of this is killing. Even though we can imagine situations in which intentionally killing one person may save the lives of many others, deontologists insist that intentional killing is always wrong. Killing is wrong even in extreme situations because it means using the person merely as a means and does not treat the human being as valuable in and of himself. Deontologists do often recognize self-defense and other special circumstances as excusing killing, but these are cases when, it is argued, the killing is not exactly intentional. (The person attacks me. I would not, otherwise, aim at harm to the person, but I have no other choice but to defend myself.)

At the heart of deontological theory is an idea about what it means to be a person, and this is connected to the idea of moral agency. Charles Fried (1978) put the point as follows:

[T]he substantive contents of the norms of right and wrong express the value of persons, of respect for personality. What we may not do to each other, the things which are wrong, are precisely those forms of personal interaction which deny to our victim the status of a freely choosing, rationally valuing, specially efficacious person, the special status of moral personality. (pp. 28–29)

According to deontologists, the utilitarians go wrong when they fix on happiness as the highest good. Deontologists point out that happiness cannot be the highest good for humans. The fact that we are rational beings, capable of reasoning about what we want to do and then deciding and acting, suggests that our end (our highest good) is something other than happiness. Humans differ from all other things in the world insofar as we have the capacity for rationality. The

behavior of other things is determined simply by laws of nature. Plants turn toward the sun because of photosynthesis. They don't think and decide which way they will turn. Physical objects fall by the law of gravity. Water boils when it reaches a certain temperature. In contrast, human beings are not entirely determined by laws of nature. We have the capacity to legislate for ourselves. We decide how we will behave. As Kant describes this, it is the difference between acting in accordance with law (plants and stones do) and acting in accordance with *the conception* of law.

The capacity for rational decision making is the most important feature of human beings. Each of us has this capacity; each of us can make choices, choices about what we will do, and what kind of persons we will become. No one else can or should make these choices for us. Moreover, we should recognize this capacity in others.

Notice that it makes good sense that our rationality is connected with morality, for we could not be moral beings at all unless we had this rational capacity. We do not think of plants or fish or dogs and cats as moral beings precisely because they do not have the capacity to reason about their actions. We are moral beings because we are rational beings, that is, because we have the capacity to give ourselves rules (laws) and follow them.

Where utilitarians note that all humans seek happiness, deontologists emphasize that humans are creatures with goals who engage in activities directed toward achieving these goals (ends), and that they use their rationality to formulate their goals and figure out what kind of life to live. In a sense, deontologists pull back from fixing on any particular value as structuring morality and instead ground morality in the capacity of each individual to organize his or her own life, make choices, and engage in activities to realize their self-chosen life plans. What morality requires is that we respect each of these beings as valuable in themselves and refrain from valuing them only insofar as they fit into our own life plans.

As mentioned before, Kant put forward what he called the *categorical imperative*. While there are several versions of it, I will focus on the second version which goes as follows: *Never treat another human being merely as a means but always as an end*. This general rule is derived from the idea that persons are moral beings because they are rational, efficacious beings. Because we each have the capacity to think and decide and act for ourselves, we should each be treated with respect, that is with recognition of this capacity.

Note the "merely" in the categorical imperative. Deontologists do not insist that we never use another person as a means to an end, only that we never "merely" use them in this way. For example, if I own a company and hire employees to work in my company, I might be thought of as using those employees as a means to my end (i.e., the success of my business). This, however, is not wrong if I promise to pay a fair wage in exchange for work and the employees agree to work for me. I thereby respect their ability to choose for themselves.

What would be wrong would be to take them as slaves and make them work for me. It would also be wrong to pay them so little that they must borrow from me and remain always in my debt. This would be exploitation. This would show disregard for the value of each person as a “freely choosing, rationally valuing, specially efficacious person.” Similarly, it would be wrong for me to lie to employees about the conditions of their work. Suppose, for example, that while working in my plant, employees will be exposed to dangerous, cancer-causing chemicals. I know this but don’t tell the employees because I am afraid they will quit. In not being forthcoming with this information, I am, in effect, manipulating the employees to serve my ends. I am not recognizing them as beings of value with their own life-plans and the capacity to choose how they will live their lives.

Case Illustration

Though utilitarianism and Kantian theory were contrasted in the case illustration about allocation of scarce medical resources, another case will clarify even more. Consider a case involving computers. Suppose a professor of sociology undertakes research on attitudes toward sex and sexual behavior among high school students. Among other things, she interviews hundreds of high school students concerning their attitudes and behavior. She knows that the students will never give her information unless she guarantees them confidentiality, so before doing the interviews, she promises each student that she alone will have access to the raw interview data, and that all publishable results will be reported in statistical form. Thus, it would be impossible to identify information from individual students.

Suppose, however, that it is now time to analyze the interview data and she realizes that it will be much easier to put the data into a computer and use the computer to do the analysis. To assure the confidentiality she promised, the professor will have to code the data so that names do not appear in the database and will have to make an effort to secure the data. She has hired graduate students to assist her and she wonders whether she should let the graduate students handle the raw data. Should she allow the graduate assistants to code and process the data?

At first glance it would seem that from a consequentialist point of view, the professor should weigh the good that will come from the research, and from doing it quickly on a computer, against the possible harm to herself and her subjects if information is leaked. The research may provide important information to people working with high school students and may help her career to prosper. Still, the advantage of doing it quickly may be slight. She must worry about the effect of a leak of information on the students. Also, since she has explicitly promised confidentiality to the student-subjects, she has to worry about the effects on her credibility as a social researcher and on social science research in general if she breaks her promise. That is, her subjects and many

others may be reluctant in the future to trust her and other social scientists if she breaks the promise and they find out.

Thus, there seem good reasons to say that from a consequentialist point of view the professor should not violate her promise of confidentiality. Fortunately, there are ways to code data before putting it into the computer or turning it over to her graduate students. She must do the coding herself and keep the key to individual names confidential.

This is how a consequentialist might analyze the situation. Interestingly, a deontologist might well come to the same conclusion though the reasoning would be quite different. The sociologist is doing a study that will advance human knowledge and, no doubt, further her career. There is nothing wrong with this as long as it does not violate the categorical imperative. The question here is whether she is treating her subjects merely as means to knowledge and her own advancement, or whether she is truly recognizing those subjects as ends in themselves. Were the sociologist to ignore her promise of confidentiality to the students, she would not be treating each subject as an end. Each student made a choice based on her pledge of confidentiality. She would be treating them merely as means if she were to break her promise when it suited her. Thus, out of respect for the subjects, the sociologist must code the data herself so as to maintain the promised confidentiality.

The two theories do not, then, come to very different conclusions in this case. However, the analysis is very different in that the reasons given for coming to the conclusion are very different. In other cases, these theories lead to dramatically different conclusions.

Our dialogue on utilitarianism and Kantian theory could continue. I have presented only the bare bones of each theory. However, in the interest of getting to the issues surrounding computers, we must move on and put a few more important concepts and theories “on the table.”

RIGHTS

So far, very little has been said about rights though we often use the language of rights when discussing moral issues. “You have no right to tell me what to do.” “I have a right to do that.” Ethicists often associate rights with deontological theories. The categorical imperative requires that each person be treated as an end in himself or herself, and it is possible to express this idea by saying that individuals have “a right to” the kind of treatment that is implied in being treated as an end. The idea that each individual must be respected as valuable in himself or herself implies that we each have rights not to be interfered with in certain ways, for example, not to be killed or enslaved, to be given freedom to make decisions about our own lives, and so on.

An important distinction that philosophers often make here is between negative rights and positive rights. Negative rights are rights that require

restraint by others. For example, my right not to be killed requires that others refrain from killing me. It does not, however, require that others take positive action to keep me alive. Positive rights, on the other hand, imply that others have a duty to do something to or for the right holder. So, if we say that I have a positive right to life, this implies not just that others must refrain from killing me, but that they must do such things as feed me if I am starving, give me medical treatment if I am sick, swim out and save me if I am drowning, and so on. As you can see, the difference between negative and positive rights is quite significant.

Positive rights are more controversial than negative rights because they have implications that are counter-intuitive. If every person has a positive right to life, this seems to imply that each and every one of us has a duty to do whatever is necessary to keep all people alive. This would seem to suggest that, among other things, it is our duty to give away any excess wealth that we have to feed and care for those who are starving or suffering from malnutrition. It also seems to imply that we have a duty to supply extraordinary life-saving treatment for all those who are dying. In response to these implications, some philosophers have argued that individuals have only negative rights.

While, as I said earlier, rights are often associated with deontological theories, it is important to note that rights can be derived from other theories as well. For example, we can argue for the recognition of a right to property on utilitarian grounds. Suppose we ask why individuals should be allowed to have private property in general and, in particular, why they should be allowed to own computer software. Utilitarians would argue for private ownership of software on grounds that much more and better software will be created if individuals are allowed to own (and then license or sell) it. Thus, they argue that individuals should have a legal right to ownership in software because of the beneficial consequences of acknowledging such a right.

Another important thing to remember about rights is the distinction between legal and moral (or natural or human) rights. Legal rights are rights that are created by law. Moral, natural, or human rights are claims independent of law. Such claims are usually embedded in a moral theory or a theory of human nature.

The utilitarian argument is an argument for creating or recognizing a legal right; it is not an argument to the effect that human beings have a natural right, for example, to own what they create. In Chapter 6 we will focus on property rights in computer software and there we will explore both natural and utilitarian property rights.

Rights and Social Contract Theories

Rights are deeply rooted in the tradition of social contract theories. In this tradition the idea of a social contract (between individuals, or between individuals and government) is hypothesized to explain and justify the obligations that

human beings have to one another. Many of these theories imagine human beings in a state of nature and then show that reason would lead individuals in such a state to agree to live according to certain rules, or to give power to a government to enforce certain rules. The depiction of a state of nature in which human beings are in a state of insecurity and uncertainty is used to suggest what human nature is like and to show that human nature necessitates government. That is, in such a state any rational human beings would agree (make a contract) to join forces with others even though this involves giving up some of their natural freedom. The agreement (the social contract) creates obligations and these are the basis of moral obligation.

An argument of this kind is made by several social contract theorists and each specifies the nature and limits of our obligations differently. One important difference, for example, is in whether morality exists prior to the social contract. Hobbes argues that there is no justice or injustice in a state of nature; humans are at war with one another and each individual must do what they must to preserve themselves. Locke, on the other hand, specifies a natural form of justice in the state of nature. Human beings have rights in the state of nature and others can treat individuals unjustly. Government is necessary to insure that natural justice is implemented properly because without government, there is no certainty that punishments will be distributed justly.

Rawlsian Justice

In 1971, John Rawls, a professor at Harvard University, introduced a new version of social contract theory (though some argue it is not a social contract theory in the traditional sense). Rawls introduced the theory in a book entitled simply *A Theory of Justice*. The theory may well be one of the most influential moral theories of the twentieth century, for not only did it generate an enormous amount of attention in the philosophical community, it influenced discussion among economists, social scientists, and public policy makers.

Rawls was primarily interested in questions of distributive justice. In the tradition of a social contract theorist, he tries to understand what sort of contract between individuals would be just. Rawls recognizes that we can't arrive at an account of justice and the fairness of social arrangements by reasoning about what rules particular individuals would agree to. He understands that individuals are self-interested and therefore will be influenced by their own experiences and their own situation when they think about fair arrangements. Thus, if some group of us were to get together in something like a state of nature (suppose a group is stranded on an island or a nuclear war occurs and only a few survive), the rules we would agree to would not necessarily be a just system. It would not necessarily exemplify justice.

The problem is that we would each want rules that would favor us. Smart people would want rules that favored intelligence. Strong people would want a system that rewarded strength. Women would not want rules that were biased

against women, and so on. The point is that there is no reason to believe that the outcome of a negotiation in which people expressed their preferences would result in rules of justice and just institutions. In this sense, Rawls believes that justice has to be blind in a certain way.

Rawls specifies, therefore, that in order to get at justice, we have to imagine that the individuals who get together to decide on the rules for society are behind a veil of ignorance. The veil of ignorance is such that individuals do not know what characteristics they will have. They do not know whether they will be male or female, black or white, highly intelligent or moderately intelligent or retarded, physically strong or in ill-health, musically talented, successful at business, indigent and so on.

At the same time, these individuals would be rational and self-interested and would know something about human nature and human psychology. In a sense, what Rawls is suggesting here is that we have to imagine *generic* human beings. They have abstract features that human beings generally have (i.e., they are rational and self-interested). And, they have background knowledge (i.e., general knowledge of how humans behave and interact and how they are affected in various ways).

According to Rawls, justice is what individuals would choose in such a situation. Notice that what he has done, in a certain sense, is eliminate bias in the original position. Once a society gets started, once particular individuals have characteristics, their views on what is fair are tainted. They cannot be objective.

So, justice, according to Rawls is what people would choose in the original position where they are rational and self-interested, informed about human nature and psychology but behind a veil of ignorance with regard to their own characteristics. Rawls argues that individuals in the original position would agree to two rules. These are the rules of justice and they are “rules of rules” in the sense that they are general principles constraining the formulation of specific rules. The rules of justice are:

1. Each person should have an equal right to the most extensive basic liberty compatible with a similar liberty for others.
2. Social and economic inequalities should be arranged so that they are both (a) reasonably expected to be to everyone’s advantage and (b) attached to positions and offices open to all.

These general principles assure that no matter where an individual ends up in the lottery of life (in which characteristics of intelligence, talents, physical abilities, and so on, are distributed), he or she would have liberty and opportunity. He or she would have a fair shot at a decent life.

While Rawls’ account of justice has met with criticism, it goes a long way in providing a framework for envisioning and critiquing just institutions. This discussion of Rawls is extremely abbreviated as were the accounts of Kant and utilitarianism. Perhaps the most important thing to keep in mind as we

proceed to the issues surrounding computer and information technology is that rights-claims and claims about justice and fairness generally presume a much more complicated set of claims. Such claims should never be accepted as primitive truths. The underlying argument and embedded assumptions should be uncovered and critically examined.

VIRTUE ETHICS

Before moving on to the ethical issues surrounding computer and information technology, one other tradition in ethical theory should be mentioned. In recent years, interest has arisen in resurrecting the tradition of virtue ethics, a tradition going all the way back to Plato and Aristotle. These ancient Greek philosophers pursued the question: What is a good person? What are the virtues associated with being a good person? For the Greeks *virtue* meant excellence, and ethics was concerned with excellences of human character. A person possessing such qualities exhibited the excellences of human good. To have these qualities is to function well as a human being.

The list of possible virtues is long and there is no general agreement on which are most important, but the possibilities include courage, benevolence, generosity, honesty, tolerance, and self-control. Virtue theorists try to identify the list of virtues and to give an account of each—What is courage? What is honesty? They also give an account of why the virtues are important.

Virtue theory seems to fill a gap left by other theories we considered, because it addresses the question of moral character, while the other theories focused primarily on action and decision making. What sort of character should we be trying to develop in ourselves and in our children. We look to moral heroes, for example, as exemplars of moral virtue. Why do we admire such people? What is it about their character and their motivation that are worthy of our admiration?

Virtue theory might be brought into the discussion of computer technology and ethics at any number of points. The most obvious is, perhaps, the discussion of professional ethics, where we want to think about the characteristics of a good computer professional. Good computer professionals will, perhaps, exhibit honesty in dealing with clients and the public. They should exhibit courage when faced with situations in which they are being pressured to do something illegal or act counter to public safety. A virtue approach would focus on these characteristics and more, emphasizing the virtues of a good computer professional.

INDIVIDUAL AND SOCIAL POLICY ETHICS

One final distinction will be helpful. In examining problems or issues, it is important to distinguish levels of analysis, in particular that between macro and

micro level issues or approaches. One can approach a problem from the point of view of social practices and public policy, or from the point of view of individual choice. Macro level problems are problems that arise for groups of people, a community, a state, a country. At this level of analysis, what is sought is a solution in the form of a law or policy that specifies how people in that group or society ought to behave, what the rules of that group ought to be. When we ask the following questions, we are asking macro level questions: Should the United States grant software creators a legal right to own software? Should software engineers be held liable for errors in the software they design? Should companies be allowed to electronically monitor their employees?

On the other hand, micro level questions focus on individuals (in the presence or absence of law or policy). Should I make a copy of this piece of software? Should I lie to my friend? Should I work on a project making military weapons? Sometimes these types of questions can be answered simply by referring to a rule established at the macro level. For example, legally I can make a back-up copy of software that I buy, but I shouldn't make a copy and give it to my friend. Other times, there may be no macro level rule or the macro level rule may be vague or an individual may think the macro level rule is unfair. In these cases, individuals must make decisions for themselves about what they ought to do.

The theories just discussed inform both approaches, but in somewhat different ways, so it is important to be clear on which type of question you are asking or answering.

CONCLUSION

While the focus of our attention will now shift to the ethical issues surrounding computer and information technology, the deep questions and general concerns of ethical theories will continue to haunt us. The dialogue is ongoing. Remember that science is never done. In both science and ethics, we look for reasons supporting the claims that we make, and we tell stories (develop arguments and theories) to answer our questions. We tell stories about why the physical world is the way it is, why human beings behave the way they do, why lying and killing are wrong, and so on. The stories we tell often get better and better over time. They get broader (more encompassing) and richer, sometimes more elegant, sometimes allowing us to see new things we never noticed before. The stories generally lead to new questions. So it is with ethics as well as science.

Computer ethics should be undertaken with this in mind, for the task of computer ethics involves working with traditional moral concepts and theories, and extending them to situations with somewhat new features. The activity brings insight into the situations arising from use of computer and information technology, and it may also bring new insights into ethical concepts and theories.

STUDY QUESTIONS

1. What are the activities involved in doing philosophical ethics?
2. How do descriptive (empirical) claims and prescriptive (normative) claims differ? Give examples of each kind of claim.
3. What is ethical relativism? What is its positive claim? What is its negative claim?
4. What are the three types of evidence often used to support ethical relativism?
5. Does this evidence support ethical relativism?
6. What are the three problems with ethical relativism?
7. What is utilitarianism?
8. What is the difference between an instrumental good and an intrinsic good?
9. Why do utilitarians believe that happiness is the ultimate basis for morality?
10. What is the difference between act-utilitarianism and rule-utilitarianism?
11. What is the traditional criticism of utilitarianism?
12. Why can't happiness be the highest good for humans according to deontologists?
13. What is the categorical imperative?
14. How can rights be based on deontological theory? How can rights be based on utility theory?
15. What are the two principles of justice according to John Rawls?
16. How does virtue ethics differ in focus from other theories discussed in this chapter?
17. What is the difference between macro and micro level issues?

SUGGESTED FURTHER READING

- BECKER, LAWRENCE and CHARLOTTE BECKER, *The Encyclopedia of Ethics* (New York: Garland Publishing, 1992).
- RACHELS, JAMES, *Elements of Moral Philosophy*. 3rd ed. (Boston: McGraw-Hill, 1999).
- SINGER, PETER, *Practical Ethics*. 2nd ed. (Cambridge, MA: Cambridge University Press, 1993).