



COMPUTER ETHICS

FOURTH EDITION

DEBORAH G. JOHNSON

CHAPTER 2

Ethics and Information Technology



CHAPTER OUTLINE

Introduction: “Doing” Ethics

- Descriptive/Normative
- The Dialectic Method
- “Ethics Is Relative”

Ethical Theories and Concepts

- Utilitarianism
- Intrinsic and Instrumental Value
- Acts versus Rules
- Critique of Utilitarianism
- Case Illustration
- Deontological Theory
- Case Illustration
- Rights
- Rights and Social Contract Theory
- Virtue Ethics
- Analogical Reasoning in Computer Ethics

Conclusion

Study Questions

INTRODUCTION: “DOING” ETHICS

In Chapter 1, we asked and answered the “why computer ethics?” question. Even though much of our attention there was on technology in general and not specifically on computers and IT, we can think of Chapter 1 as addressing the “computer” part of computer ethics. This chapter will address the “ethics” part.

The meaning of the term “ethics” is not easy to specify, and yet much of the controversy and skepticism about ethics seems to arise from ill-conceived notions of ethics. For example, some think of ethics as a set of rules that are universally binding on all people at all times; as such, they presume that moral rules must be derived from some sort of transcendental or higher authority, such as God or human nature or reason. Using this conception of ethics, it is easy to become skeptical when most attempts to identify the universally binding rules fail. Yet there are a variety of alternative approaches to ethics, approaches that have little to do with universal rules. Some emphasize right and wrong actions, others emphasize good and bad consequences, yet others emphasize virtue or justice. Moreover, a distinction should be made between theoretical and practical ethics. Theoretical ethics focuses on giving an account of morality, what it is, and how its claims are justified. Practical ethics draws on theoretical ethics, but is grounded in social practices and aims to understand, illuminate, and develop strategies for practical decision making.

The approach taken here is practical. Ethics is understood here to refer to a way of looking at human conditions and interactions using a set of concepts and theories that are distinctively normative. Ethics is a normative lens through which to view human arrangements, choices, and actions. [The meaning of the term “normative” will be explained further in a moment.]

We live in a world that requires making choices and acting. Our deliberations about how to act and what to choose often involve moral notions (right and wrong, loyalty, duty, justice, responsibility), ethical principles (do no harm, tell the truth, keep your promises), and ideas about what makes for a full and meaningful life (concern for others, community, friendship). In this book, we use analytical methods to illuminate the ethical aspects of situations and the ethical implications of deciding one way or another, or adopting one policy or another. In this way, we are framing computer ethics as a form of practical ethics. Although we draw on theoretical ethics, our primary interest is the ethical analysis of real situations, situations in which IT plays a role. As already mentioned, theoretical ethics is concerned with explaining the very idea of morality and with understanding the foundation of moral claims, typically universal moral claims. Our more modest goal here is to provide analysis that informs (although does not necessarily dictate) decision and action. The framework, concepts, and theories discussed here are intended to help readers think through situations that arise in the real world and reflect on what a better world would look like.

Sometimes the lens of ethics brings to light an aspect of a situation that seems to trump all other aspects. For example, suppose you are contemplating lying about a product that you are selling—say the product is a toy and you know that the paint on the toy contains a dangerous amount of lead. Here it would seem that no matter how you look at the situation—from the economic, legal, or cultural perspective—lying

about the product seems wrong. Here ethics “trumps” all other aspects of the situation. Although this sometimes happens, not all situations are so clear. Often the ethical implications are intertwined with other dimensions of life—legal, economic, religious, political. Thus, the definition of ethics we will use here does not presuppose a priority to ethics. It does, nevertheless, presume that deliberation and action are better when the ethical aspects of a situation are taken into account.

The concepts and theories explained in this chapter come from a long history of philosophical thought. Philosophers have developed theories that explain the idea of morality, and have argued for various systems of ethical decision making. As with all fields, however, the state of understanding continues to change, with current ideas being contested, new ones offered, and the body of knowledge growing. This chapter scratches only the surface of a complex, heterogeneous, evolving body of knowledge. Although this chapter provides a quick sketch of ethical concepts and theories, our aim is to jumpstart a dialogue on ethics that readers will continue in academic courses and throughout their lives, in their personal reflection, and in ongoing conversation with others.

Descriptive/Normative

The study of ethics is normative. When individuals or groups make decisions and act, the more they know about the state of the world, the better. However, having an accurate description of states of affairs in the world is only part of what is involved in acting wisely. Decisions and actions are aimed at the future. They are normative in the sense that they can go one way or another, and one chooses a direction when one acts. When one acts, one says, in effect, “this is what I want to happen” or “telling the truth will be better than lying” or “buying this television will make me happy” or “voting for Smith is more likely to lead to improvements in the city.” These statements are all normative; they implicitly have to do with what is good/bad or better/worse or worthy/unworthy. So, although it is true that the more one understands the world in which one acts, the better decisions are likely to be, no matter how accurate one’s understanding of the world, one ultimately has to make choices, and choices involve much more than the way the world is. Ethics has to do with steering one’s life, making intentional choices, and contributing to the future. If you want to avoid doing harm and contribute to the improvement of human conditions, it is essential that you think about what constitutes a better condition, what makes for a better, more just, more peaceful, and more fulfilling world. In this respect, ethics is about ends. Moral actions, rules, principles, or guidelines are all aimed at achieving ends.

The distinction between descriptive and normative claims is important here although it is, by no means, simple. *Descriptive* statements are statements that describe a state of affairs in the world. For example: “The car is in the driveway”; “Georgia is south of Tennessee”; “XX percent of Chinese citizens have Internet access in their homes”; “XX percentage of traffic on the Web is to pornographic websites.” These are all *empirical* claims in the sense that they can be verified or proven false by looking and seeing. Observations can be made, surveys can be administered, and individuals can be asked, although this isn’t always easy to do. Consider the difficulties of verifying the following claim: “All societies consider some domain of life private,

although which domain(s) of life is considered private varies a good deal from society to society.” Verifying this claim would involve not only examining all societies, it would involve clarification as to what it means for a society to consider an area of life private. Nevertheless, the claim is descriptive; it is a claim about conditions in the world, conditions that can be examined to see whether the claim is accurate.

By contrast, normative claims are prescriptive and evaluative. Keeping with the above example, someone might claim: “Every society *should* keep some domains of life private.” This is not an empirical claim; it cannot be verified by examining societies. The claim makes a recommendation and although empirical evidence might be brought in to support the claim, ultimately what *is* the case and what *ought* to be the case are different matters.

Social scientists gather empirical data and report their findings on a wide range of topics including moral attitudes and behavior. For example, psychologists and sociologists might identify the processes by which children develop moral concepts and sensibilities. Or they might measure how individuals value and prioritize various goods such as friendship, privacy, and autonomy. When anthropologists study a culture, they describe complex moral rules in the culture they observe. They are describing lived and observed moral systems. Similarly, historians may trace the development of a particular moral notion in an historical period. These historical and social scientific studies are descriptive; they examine morality as an empirical phenomenon. They do not, however, tell us what is right and wrong. They don’t tell us what people *should* do, only what people, in fact, do. On the other hand, normative analysis deals with prescriptive and evaluative claims.

Earlier we said that the approach taken here is aimed at a kind of analysis that would be helpful in decision making and acting. That is precisely what normative analysis does. It is concerned with evaluating and critiquing states of affairs in search of ways to think about what was wrong or what would be better, a better state of affairs, better social arrangements, a better way to treat one another, ultimately to inform action. Ethical analysis is directed at human ends and goals—how we should treat one another, what constitutes justice and fairness, what we owe one another in virtue of being human, and when we should restrain our personal interests and desires. Making decisions, choosing, and setting policies are all intrinsically normative endeavors.

Normative claims cannot be supported simply by pointing to the facts about what individuals do or say or believe. Likewise, descriptive issues cannot be resolved by claims about what ought to be or what is just and fair. For example, although it is descriptively accurate to say that throughout human history some individuals have intentionally killed others, you probably wouldn’t infer from this that it is okay for individuals to kill others when they choose (a normative claim). On the other hand, it is not uncommon to hear individuals justify the downloading of proprietary music on grounds that it is commonly done (even though it is illegal). Here there is what seems to be an invalid inference from a descriptive claim—“it is commonly done”—to a normative claim—“it’s okay for me to do it.” When we reflect on the reasoning here, it is difficult to see how the *descriptive* claim justifies the *normative* claim. The fact that individuals often engage in illegal behavior doesn’t seem to tell us anything about whether the behavior is right or wrong.

On the other hand, the two kinds of claims often can be productively played off one another. Empirical information may be helpful in identifying ways of thinking about a normative issue. For example, exploring why individuals believe that downloading music is okay may provide some ideas that help to identify normative arguments or the moral principles at issue. Moreover, normative beliefs often influence which and what kind of empirical data we collect. For example, social scientists seek information about the degree to which citizens of various countries are using the Internet because they believe (normatively) that the spread of the Internet is an extremely important social phenomenon (that it is important for economic development, the spread of democracy, etc.).

Thus, although the goal in this book is to generate normative insights and analysis, we will use descriptive claims and evidence when it is helpful to do so. We will never, however, use a descriptive claim as the primary justification for a normative claim.

The Dialectic Method

How, you might now ask, does one “do” ethics? When it comes to describing moral beliefs and practices, we examine what people think and do, and gather and reflect on empirical information. However, facts and descriptions are not enough. Normative analysis generally involves identifying a principle or value, exploring what the principle or value implies, and making a case for a position. In practical ethics, this means connecting the principle or value to a particular situation, and considering arguments for various courses of action or decisions with regard to the situation. For example, in the virtual rape case described in Chapter 1, we might begin by trying to identify the behavior in question and link it to a moral concept. Rape is wrong but did Bungle or the person behind Bungle commit rape? If not, then what was the wrong? Can we think of the behavior as a violation of an implicit community standard? If so, then we would have to explain why community standards are so important. Or we might link the behavior to the harm associated with exposing individuals to sex and violence without warning them. If neither of these strategies work, then we have to find another way to characterize the behavior that connects it to a moral norm or principle.

Once a value or principle has been identified, ethical analysis proceeds with what is often referred to as a dialectic process. Here it is important to note that consistency and coherence are important tools for analysis. Using the dialectic method, normative claims are formulated into arguments. An argument is simply a claim and a set of reasons that justify the claim. Once arguments are formulated, they can be examined for their coherence, plausibility, and consistency, as well as for their fit with ordinary experience and relevant empirical information.

To understand the dialectic method, consider your own experience with discussions of ethical issues. You have probably witnessed, if not participated in, heated debates about euthanasia, abortion, affirmative action, and the distribution of wealth. Or consider discussions about downloading proprietary music, government surveillance of e-mail, or using robots to take care of the elderly. Often when individuals are asked to explain why they think a type of behavior or a policy is wrong, they have

difficulty articulating their reasons. The first step in the dialectic process is to move from unreflective beliefs and gut feelings to claims that are connected to a value or principle that others are likely to accept. Unexamined claims can be the starting place for ethical analysis, but they are only starting places. Using the dialectic method, the reasons the individual has for making a claim have to be “put on the table.” Why, we have to ask, would anyone claim that censorship is wrong, that downloading music isn’t stealing, or that relying on robots to make decisions is dehumanizing?

If reasons for a moral belief cannot be put forward, then there can be no dialogue. More importantly, if an individual cannot give reasons for his or her moral beliefs or opinions, then it would seem there is nothing to recommend them. If I don’t understand why you believe what you do, I have no “reason” to believe what you believe.

Discussions of ethical issues that stay at the level of statements of belief without reasons tend to end quickly with statements like “everyone is entitled to his or her own opinion.” There is little point in talking about ethics in this way, except perhaps to see where others stand. The dialectic method proceeds by insisting that we each give reasons for our moral beliefs so that the reasons can be examined and critically evaluated.

The critical evaluation is often done in the context of trying to convince someone to reject a position, or to adopt another position, but it can 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 the dialectic method, 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, the dialectic 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. If the claims appear to be inconsistent, then you should either change one of the claims or provide an account of how the two seemingly disparate positions are, in fact, consistent. 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 the dialectic method, consider first a case that does not involve IT. Suppose you start out by making the claim that euthanasia is wrong. You articulate a principle as the reason for this claim. Say, for example, 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 severely brain damaged? When the person is terminally ill? When the person is young or elderly? Because your principle concerns the value of human life, it has implications beyond the issue of euthanasia. 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 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, although it is still not permissible to intentionally kill, it is morally permissible to let a person die.

Whether the dialogue is inside your head (your own personal reflection), or a discussion with others, as it progresses, it leads to a more and more precise specification of the claim and its defense. 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. Nevertheless, the dialectic does not always lead to a final and absolute conclusion. Nor will the dialogue necessarily lead to unanimous agreement among the discussants. Good dialecticians are always open to further discussion with the idea that even if you don't change your mind, every discussion is an opportunity to learn more, see another connection or aspect, and to hear another perspective.

We can illustrate the dialectic method further with a situation involving IT. Consider the following case described in a recent law journal article:

On September 7, 2005, a former Delta Air Lines flight attendant filed a federal sexual discrimination lawsuit claiming that she was suspended and later fired because of material she posted on her personal blog. Ellen Simonetti was laid off after her "Queen of the Sky" blog showed a picture of her in her Delta uniform. The blog, a moderately fictionalized account of life in the air, never named Delta as her employer, but one photo did show a pin indicating she worked for the airline. Delta's decision to terminate her was based on "inappropriate photographs" of plaintiff in her uniform on the website. Ms. Simonetti claims that she was not aware of any company anti-blogging policy. According to a BBC News source, "there is guidance which suggests the company uniform cannot be used without approval from management, but use in personal pictures on websites is unclear."

[T. Watson and E. Piro, "Bloggers beware: a cautionary tale of blogging and the doctrine of at-will employment" *Hofstra Labor & Employment Law Journal*, Winter, 2007]

In this case, the first step of linking the situation to a moral concept or theory may seem easy. What is at issue here is freedom of expression and censorship. The company (Delta) seems to want to prevent one of its employees from speaking freely (posting information); hence, it seems to be a case of interfering with freedom of expression. The company wants to censor its employee's blog. Once the case is linked to a concept, we have to explore the fit and see whether the concept can be used to illuminate, or even decide, the case. Is it a cut-and-dried case of an employer interfering with an

employee's right to freedom of expression? Following this path, we would want to explore further the defense of freedom of expression. Do individuals have a right to freedom of expression? Why? Such a right might be defended by referring to legal rights. In the United States, this argument might be framed around the First Amendment. However, an argument might also be made that freedom of expression is a natural or human or moral right. Arguments have to be formulated and examined. Whatever way one goes on rights, the question that will come up is, "is the right absolute, or are there situations in which restrictions on free speech are justified?" One commonly noted case is that no one is allowed to yell "fire" in a crowded place. And, there are other domains in which free speech is restricted. Hate speech is a case in point.

One way or another, the dialectic is likely to move in the direction of employer–employee rights because employers do have the right to require their employees to sign agreements to keep trade secrets confidential, and they have the right to protect their reputation. Moreover, in U.S. law, employer–employee relationships are covered by what is called the "doctrine of at-will employment," which means that employers can fire employees with little cause. By moving the dialogue in this direction, the case is reframed as one that raises questions about the boundaries of employer–employee rights and obligations. In the law article from which the case is taken, the authors note that blogs are a new phenomenon so there are no legal precedents, but they seem skeptical that employee-bloggers will be protected from being fired for what they post on their blogs.

The dialogue can go in any number of directions, and our cursory look at euthanasia and blogging merely suggests how moral concepts and principles come into play and are then used to understand a situation and develop arguments.

As mentioned earlier, the dialectic method does not always lead to a definitive conclusion about what should be done or what precisely was wrong, but it almost always leads to better understanding. Thus, it is important to keep in mind at the onset that understanding can be improved and progress made, even when one has not reached absolute conclusions. Through the dialectic we learn which arguments are weaker and stronger, and why. We come to better understand the ideas that underpin our moral beliefs. We develop deeper and more consistent beliefs, and come to understand how moral ideas are interrelated and interdependent. The dialectic and the analysis show us what is at stake, help us to understand the values and interests relevant to various actors, and often help us to identify alternative forms of action or decision making.

When it comes to practical ethics, there seems no reason to believe that there is, or has to be, a single right answer to an ethical problem. Whitbeck (1998) argues that ethical problems are better understood on the model of design problems. When you give a design problem to multiple teams of engineers specifying the design requirements, you generally get different designs from each team. Even when you specify the features you want all designs to meet, engineers will creatively balance various factors against one another. Suppose you ask teams to design a car seat that meets regulatory requirements for safety, weighs no more than a specified amount, and costs no more than a certain amount to manufacture. Each team will come up with a different design, that is, using different materials, having a different shape, and with differing

accessories. If we understand ethical decision making on the design model, there is no reason to believe that there is only one right way to act when one finds oneself in an ethical dilemma. On the other hand, it is important to note that thinking of ethical problems on the model of design problems does not lead to a free-for-all or anything-goes in ethics. As Whitbeck explains:

Although no unique correct solution may exist, nonetheless, some possible responses are clearly unacceptable—there are wrong answers even if there is no unique right answer—and some solutions are better than others.

So it is with ethical issues. We may rule out some solutions to an ethical dilemma as utterly unacceptable. We may find a range of possible courses of action with varying advantages and disadvantages. We may not be able to identify a single action that is “the” right one or “the” morally obligatory one, and yet we can still conclude that we must do something. The dialectic process helps to sort out what actions are entirely unacceptable, and distinguish possible courses of action with various advantages and disadvantages.

As you will see in a moment, a familiarity with traditional moral concepts and theories will help in linking situations to moral concepts and theories and formulating reasons and arguments. Ethical theories provide frameworks in which arguments can be cast. Moreover, ethical theories provide common ground for discussion. They establish a common vocabulary and frameworks within which, or against which, ideas can be articulated. However, before introducing these concepts and theories, it will be helpful to further illustrate the dialectic method while exploring a notion that may come into play as you begin to think about and discuss ethical issues.

“Ethics Is Relative”

Many discussions of ethics begin with someone putting on the table the idea that “ethics is relative.” Ethical beliefs depend, they claim, on what country you live in, where you were born, your age, or your personality. Claims of this type are also sometimes used to end debates about delicate issues such as abortion or euthanasia. That is, someone may conclude the discussion by saying: “everyone is entitled to his or her own opinion” or “I guess right and wrong depend on where you are sitting.” Although seemingly simple, when subjected to the scrutiny of the dialectic method, claims of this kind turn out to be quite complex and perhaps confused. Hence, the “ethics is relative” claim is a good starting place to further illustrate the dialectic method.

To better understand what someone might have in mind when he or she claims that “ethics is relative,” we can begin by using the descriptive–normative distinction drawn earlier. Is “ethics is relative” a descriptive or normative claim? What sort of justification might be given in each case? If “ethics is relative” is taken to be a descriptive claim, that is, a claim about what people think and do, then we can reformulate it into the following, more specific, claim: “ethical beliefs, rules, and practices vary from culture to culture and from time to time.” Indeed, if this is what “ethics is relative”

means, a good deal of evidence can be put forward to support it. Typically, three kinds of evidence are put forward in support:

1. At any given time (including the present), there is a great deal of variation in what particular individuals and groups consider right and wrong. For example, it is considered immoral for women to appear in public without their faces covered in some societies; what some consider to be bribery is common practice in certain places, an ordinary part of doing business; and polygamy is permissible in some cultures.
2. Moral norms vary over time so that what was considered wrong at one time, in a given society, may be considered right at another time. Slavery is a good example, as well as prohibitions on sex before marriage, and the use of physical force to punish children. The moral status of such practices has changed over time.
3. Moral beliefs seem to be largely influenced by when, where, how, and by whom one is raised. If I had been born in certain parts of the world, 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.

All three of these types of evidence are empirical claims that can be, and have been, substantiated by historians, anthropologists, sociologists, and psychologists. Although one might argue with a specific detail, in general the evidence seems strong, perhaps even undeniable. When “ethics is relative” is understood to be a descriptive claim, it seems to be highly plausible. Nevertheless, there are a number of ways we can put the claim to further testing through the dialectic method. One way is to press what the claim means or implies; another is to reformulate “ethics is relative” into a normative claim, and then see whether it is defensible as such.

Taking the first tack, we can press deeper, not by attacking the evidence but by questioning whether it supports the conclusion. In other words, one might wonder whether the diversity of belief shown by the evidence isn’t superficial and misleading. Isn’t it possible that universal norms underlie the seemingly disparate beliefs, rules, and practices? Universal norms may be at work in all human societies, but hidden from sight because they are expressed or interpreted in different ways in different contexts. General moral norms such as respect for human life or maximizing happiness might be operative, even though these general norms get expressed in different ways, at different times, and in different places. Anthropologists often draw attention to a seeming universal prohibition on incest although, of course, societies have very different ideas about kinship (i.e., which persons are forbidden as sexual partners). The point is that even if the “ethics is relative” claim is taken to be essentially a descriptive claim, it can be challenged and the dialectic method used to follow out the challenge. Are there any universal norms? How do we explain when individuals defy and rebel against their society’s norms? How do moral norms change?

A second way to move the dialectic forward is to treat “ethics is relative” as a normative claim. In some sense, when we took the claim to be descriptive, it didn’t seem to be a moral claim at all. That is, it didn’t provide any recommendations or guidance as to how we “ought” to behave; it didn’t give us a rule or principle that we have to use in

making decisions. One way to reformulate “ethics is relative” into a normative claim is to interpret it to mean that “right and wrong *are* relative,” that is, whatever is right or wrong is nonuniversal and depends on something like one’s culture or when and where one is living. Here the claim might be comparable to: “what is right for you may not be right for me” or “when in Rome, do as the Romans.” Sometimes ethical relativists (i.e., those who claim “ethics is relative”) seem to assert that right and wrong are relative to the individual, and other times that right and wrong are relative to the society in which one lives. Each formulation would take the dialectic in a different direction.

Pursuing the latter alternative, “ethics is relative” would mean that what is morally right for me, an American living in the twenty-first century, differs from what is right for a person living in another country or in another time period. In other words, the claim seems to be that right and wrong are relative to one’s society, and that one *should* act in conformance with the rules of one’s society.

So our dialectic has led to a clear formulation of the claim, clear enough for it to be tested in the dialectic. When we turn a critical eye to this claim, it appears to be quite problematic. The claim that “one should act in conformance with the rules of one’s society” runs into at least three serious problems.

First, although normative ethical relativists have a variety of ways to articulate and defend their claim, some versions of ethical relativism seem to slip into inconsistency and even self-contradiction. If normative ethical relativists say that right and wrong are relative to one’s society, and mean by this that an individual is bound by the rules of his or her society and should follow the rules in their society, then the relativist seems to be asserting a universal moral principle. “Everyone,” they claim, “ought to follow the norms of their society.” So, if this is what relativists mean, they contradict the very claim they make: it is contradictory to say that ethics is relative *and* “everyone” ought to follow the same general principle. To be sure, ethical relativists can try to defend against this criticism, but notice that if they pull back from making any normative claim whatsoever, then it would seem they don’t have an ethical theory at all, but merely a description of the variation in moral beliefs and practices.

Another potential inconsistency arises when one considers a common motive for making the relativistic claim. Some ethical relativists adopt ethical relativism because they are trying to stop what anthropologists call “ethnocentrism.” Ethnocentrism refers to people from one culture using the standards of their own culture to judge (and likely condemn) the practices and people of another culture. Avoiding ethnocentrism means being tolerant and respectful of difference. It means appreciating the diversity and variety of beliefs and practices, including moral beliefs and practices. However, this stance against ethnocentrism is not exactly consistent with ethical relativism, at least not with normative ethical relativism. If you adopt the position that it is wrong to judge other cultures by the standards of your own, then you seem to be asserting a universal moral principle. You seem to be affirming one rule that goes beyond, or has special status beyond, all the others that you take to be relative. So, once again, it seems that there is something contradictory about normative ethical relativism, at least the versions we have considered.

Second, if the normative ethical relativist claim is that we ought to follow the rules of our society, then what is being claimed is antithetical to a good deal of

human experience with moral heroes. Many of our most highly regarded moral heroes—Socrates, Martin Luther King, Ghandi, even Jesus, would, on this account, be considered wrong or bad because they did not follow the rules of their society. Adopting the normative claim that one should follow the rules of one’s society seems to rule out resistance or rebellion in situations that are often considered tests of moral virtue.

Finally, the “ethics is relative” claim does not provide much help in making moral decisions, especially not with decisions in tough situations. Many ethical decisions are easy to make; we know we should keep our promises, avoid intentionally harming others, refrain from stealing, and so on. We tend to look to moral principles and theories when cultural practices are unclear or novel situations arise. Many of the most daunting and important ethical issues individuals and societies face are those arising from new technologies that create situations that humans haven’t faced before—should I donate my organs for transplantation? Should we allow human cloning? Should ISPs filter child pornography and prevent their users from accessing it? These are precisely the kinds of ethical questions that cannot be decided by social convention because there are no absolute rules and practices that precisely apply. Thus, a moral principle that says right and wrong are relative or “you ought to do what is considered right in your society” just isn’t very helpful.

Thus, although many other moves can be made in a dialectic about ethics, the versions of “ethics is relative” that we have considered do not seem plausible. You can now take the dialectic in another direction.

Because we have not drawn a definitive conclusion, it is important to point out that we have made progress. We have clarified the claim that “ethics is relative” by distinguishing a descriptive interpretation and a normative interpretation. We have examined evidence put forward to support the descriptive claim, and have evaluated the evidence. We have identified three problems with a normative interpretation of “ethics is relative.” As a normative claim, it seems to be self-contradictory; it seems inconsistent with our ideas about moral heroes, and doesn’t seem to provide the kind of guidance we often seek from ethical theories. Most importantly, we have illustrated the dialectic method that we will continue to use throughout this book.

ETHICAL THEORIES AND CONCEPTS

We turn now to several moral theories that have stood the test of time in moral philosophy. They provide frameworks and vocabulary for engaging in the dialectic process, although they are themselves subject to the scrutiny of the dialectic method. None of these theories is meant to provide an algorithm for ethical decision making; rather they provide modes of thinking, tools to use in analyzing ethical issues.

Utilitarianism

Utilitarianism is an ethical theory claiming that what makes behavior right or wrong depends wholly on the consequences. For this reason it is also often classified as a form of “consequentialism.” 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. In one version of utilitarianism, what is all important is happiness-producing consequences. 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 good consequences.

According to the version of utilitarianism that we will use, individuals should adhere to a basic principle: *Everyone ought to act so as to bring about the greatest amount of happiness for the greatest number of people.* Following the dialectic method, we should now ask: What, if any, support can be given for this theory? Why should we act to bring about the greatest amount of happiness? Why shouldn't we each seek our own interest? Utilitarianism has an answer.

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 that human beings seem to value, 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 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 animals 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, because something that 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.

The version of utilitarianism on which we are focusing 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 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 stop only 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.

Utilitarians claim that any discussion of what you should seek in life, and what is valuable, will not stop until you get to happiness. 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 always center on whether or not one has chosen the correct means to happiness. The value of happiness isn't questioned because happiness is intrinsically good.

So, when a person is faced with a decision about what to do, the person should consider possible courses of action, predict the consequences of each alternative, and choose that action which brings about the most good consequences, that is, the most happiness. The utilitarian principle provides a rough decision procedure. When you are choosing between courses of action, the right action is the one that produces the most overall net happiness (happiness minus unhappiness). 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 one's self. 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 will, on balance, bring about more good (consequences) than harm (bad consequences). At the core, cost-benefit or risk-benefit analyses are utilitarian. Benefits are weighed against risks. For example, if a community were considering whether to allow a new waste

disposal plant to be built in their area, the community would weigh the benefits of having the plant there against the risk of harm and other negative consequences to all those who would be affected.

Acts versus Rules

Because of disagreements on important details, philosophers have formulated different versions of utilitarianism. One important and controversial issue 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 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 try 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 is 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, utilitarians propose that the utilitarian principle be used to decide the laws of a society. Or they point out that the laws we currently have can be justified on utilitarian grounds. Prohibitions on stealing, killing, breaking contracts, and fraud, for example, are justified because of their consequences for human well being. 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, because 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.

Although we cannot pursue the link here, it is worth noting that utilitarianism might be used to return to our earlier discussion of “ethics is relative” because utilitarianism might be thought of as capturing part of the idea of relativism. Because the theory claims that the right thing to do depends on the consequences *and* because the same action performed in one context or set of circumstances may produce quite different consequences in another context, utilitarianism seems to allow that the right thing will vary with the context. For example, although in general more good may result from telling the truth, lying may be better in certain circumstances. Even rule-utilitarians must admit that the rules that will produce the most happiness may vary from situation to situation. A simple example would be to imagine that in a natural environment in which water is scarce, a moral prohibition on using water in swimming pools or to water lawns would be justified. 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 is compatible with varying laws and moral practices at different times or in different places.

Now that the fundamentals of utilitarianism have been explained, it is worth remembering, once again, that we are engaged in a dialectic process. We have described the idea of utilitarianism and have made a 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 can be put to the test of critical evaluation.

Critique of Utilitarianism

One of the most 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, because 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 to be justifiable (if not obligatory) in a utilitarian framework. Another more contemporary example is to imagine a situation in which by killing one person and using all his or her 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 because 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, utilitarians can argue that utilitarianism does not justify such unsavory practices. Critics, they may 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 practices such 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 they happen to be at the hospital at a time when a major accident occurred and organs were needed to save many victims. Thus, the good effects of practices of this kind could never counterbalance the long-term 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. The current debate about the use of torture to extract information that might prevent events such as terrorist attacks fits this form of analysis. Although most agree that torture is bad, some argue that the bad is counterbalanced by the good consequences that may result. Others argue that in the long-run it does more harm even to the torturing country because it means that their soldiers are more likely to be tortured if torture becomes a common practice. Still others argue that tortured prisoners are likely to say anything at all during torture, and that makes intelligence gathered during torture largely useless because the truth must still be sifted out from the lies.

In dialectic analysis, it is important to pick up on our strongly held moral intuitions because 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) hints at something missing in utilitarianism

and points in an alternative direction. A concrete case will help us further understand utilitarianism and introduce an alternative 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 to decide 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 distribution method implied a principle that is antithetical to utilitarianism. It suggested that each and every person, no matter what their 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 a means to the betterment of society, and by that criteria some individuals are much more valuable than others.

In an ideal world, more kidney dialysis machines would be produced so that no one would have to do without. At the time, this was impossible (as it is now for other types of medical treatment). Because decisions had to be made, the critics of distributing access to kidney dialysis machines on the basis of social utility proposed that access should be distributed by means of a lottery that included all of those in need. In a lottery, everyone has an equal chance; everyone counts the same. This, they argued, was the only fair method of distribution.

The unfairness of the utilitarian distribution is important because it goes to the heart of the theory. Oddly, although the theory treats each individual's happiness as equal, when overall or net happiness is determined by adding up and balancing bad against good consequences, some individual's unhappiness turns out to be dispensable for the sake of the happiness of others. Critics argue that people are valuable in themselves, not for their contribution to overall happiness. They argue that utilitarianism leads to imposing an unfair burden on some individuals; it treats some individuals as means to the good of others.

Before we explore an alternative to utilitarianism, we should note that utilitarianism goes a long way in providing a reasoned and comprehensive account of many of our moral notions. It is not a theory to be dismissed lightly. Consequences seem an important element in moral reasoning and in moral practices. However, we turn now to an ethical theory that articulates the reasoning underlying the critique of utilitarianism.

Deontological Theory

In utilitarianism, what makes an action right or wrong is outside the action; it is the consequences, effects, or results of the action. 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 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 either I fear getting caught or believe I will be rewarded for doing so, then my act is not morally worthy.

We will focus here on the theory of Immanuel Kant. Referring back to the allocation of dialysis machines, Kant's moral theory justifies distribution by a lottery, or at least *not* by social value. In Kant's philosophy, one must always act according to the categorical imperative. The categorical imperative specifies that we should *never treat human beings merely as means to an end. We should always treat human beings as ends in themselves.* Utilitarianism is criticized because it appears to tolerate sacrificing some people for the sake of others. In utilitarianism, right and wrong depend 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 victim merely as a means and does not treat the human being as valuable in and of him- or herself.

Yes, deontologists recognize self-defense and other special circumstances as sometimes excusing killing, but these are cases when, it is argued, the killing isn't

intentional. In their defense someone might say: "The person attacked me; I had no other choice but to defend myself. After all, I too am of value; I did not aim at the attacker's death, I aimed only to stop the attack."

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 have the capacity to legislate for themselves. We decide how we will behave. As Kant describes this, it is the difference between behavior that is determined by the laws of nature, that is, acting in accordance with law (as plants and stones do) *and* acting in accordance with *the conception* of law. Only human beings are capable of the latter.

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 respect this capacity in others.

Notice that it makes good sense that our rationality is connected with morality, because 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 full capacity to reason about their actions. We are moral beings because we have the capacity to give ourselves rules (laws) and follow them. (Some may dispute that dogs and cats are without rational capacity; they may claim that dogs and cats and other nonhumans have the ability to conform to a conception of law. The dialectic could go off here in the direction of considering whether the rational capacity required for morality can come in degrees.)

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

¹The term "deontology" derives from the Greek words *deon* (duty) and *logos* (science). Etymologically, then, deontology means the science of duty.

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. In other words, morality requires that we don't treat others merely as a means to our own ends.

As mentioned before, Kant's moral philosophy centers around what he called the *categorical imperative*. Although he puts forward three versions of it, we have focused on the second version: *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 in ways that recognize this capacity. This is precisely what it means to respect a person.

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. For example, if I own a company and hire employees to work in my company, I might be thought of as using my employees as a means to my end (i.e., the success of my business). This, however, is not a violation of the categorical imperative so long as I treat the employees as ends in themselves, which involves paying them a fair wage, being honest about the dangers of the work environment, evaluating their work fairly, and so on. In these ways I respect my employees' abilities to choose for themselves whether they want to work for me and under what conditions. What would be wrong would be to take them as slaves and coerce them to 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

Although utilitarianism and Kantian theory were contrasted in the case illustration about the allocation of scarce medical resources, another case will clarify this 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 correlate information in the study with particular students.

Suppose, however, that it is now time to code the interview data, and she realizes that it will be much easier to have graduate student assistants do this rather than doing it herself. 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? In a utilitarian analysis, the professor would weigh the good consequences that will come from the research (and especially from getting the results out quickly) versus the possible harm to her subjects and herself if the graduate students leak information about individual students. The research will provide important information to people working with high school students and may help the professor's career to prosper. Still, she has explicitly promised confidentiality to the student-subjects and 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. Her subjects, and many others, may be reluctant in the future to trust her and other social scientists if she breaks the promise and information on individual interviewees leaks out. Moreover, the benefits of getting the research done quickly may be marginal.

From a utilitarian perspective, then, it would seem that the professor should not violate her promise of confidentiality. Fortunately, there are ways to code data before graduate students handle it. As well, there are many steps she can take to ensure that the graduate students are well informed about the confidentiality of the data and the consequences of their leaking information about individuals.

Interestingly, a deontologist is likely to come to the same conclusion, although the reasoning would be quite different. On a deontological analysis, the important question is not whether good and bad consequences will result from assuring the confidentiality of the data, but whether the professor treats her subjects merely as means to her end of developing new knowledge and advancing her own career. Is she recognizing the student-subjects as ends in themselves? Clearly, were she to ignore her promise of confidentiality to the students, she would not be treating them as ends. Each student decided for him- or herself whether to participate in the study, and each made his or her choice based on the professor's 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 ensure the confidentiality of the data and either handle the raw data herself, or put procedures in place that will ensure that graduate students keep what they see confidential. Indeed, they should be told that the consequences of revealing confidential data will be severe.

The two theories do not, then, come to very different conclusions in this case. However, the analysis is different, that is, the reasons for keeping the data confidential are distinctive. Thus, it isn't hard to imagine that the theories lead to dramatically different conclusions in other cases.

Only the bare bones of each theory have been presented. The dialectic could go off in any number of directions here. However, in the interest of getting to the issues surrounding computers and information technology, we must move on and put a few more important ideas "on the table."

Rights

So far, very little has been said about rights, although we often use the language of rights when discussing moral issues. “You have no right to say that to me.” “My boss has no right to tell me what to do on the weekends.” Ethicists often associate rights with deontological theories. The categorical imperative requires that each person be treated as an end in him- 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 him- or herself implies certain rights, for example, a right not to be killed or enslaved, a right to be told whether we are going to be used in research, a right to make decisions about how we will live our 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 counterintuitive. 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.

Although, 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 software. As we will see in Chapter 5, utilitarians argue for proprietary rights in software on grounds that much more, and better, software will be created if the individuals who create it are allowed to own (and then license or sell) it. Thus, they argue that individuals should have a legal right to ownership of software because of the beneficial consequences of creating such a right.

So, rights can be grounded in different theoretical frameworks. Distinctions also have to be made between legal, moral, natural, and human rights. Legal rights are rights created by law. Moral, natural, or human rights are claims independent of law and grounded in theories that pertain to morality, nature, or what it means to be a human being, respectively. The important point to remember is that whenever an argument is framed in terms of rights, it is a good idea to identify what kind of right is being claimed, and what theory underlies the rights-claim.

Rights and Social Contract Theory

The idea that individuals have fundamental “rights” is deeply rooted in social contract theory. In this tradition, 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. Theorists depict the state of nature (without government or civil society) as a state of insecurity and uncertainty. Thomas Hobbes, for example, describes the state of nature as “solitary, poor, nasty, brutish, and short” (*The Leviathan*). The state of nature is so miserable that rational human beings would agree (make a contract) to join forces with others and give up some of their natural freedom in exchange for the benefits of cooperation. They would agree to abide by rules and refrain from certain actions in exchange for others doing the same.

Arguments of this kind are made by several social contract theorists and each specifies the nature and limits of the obligations incurred 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. John 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.

In 1971, John Rawls introduced a new version of social contract 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, because 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. He recognizes that we cannot 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 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 just.

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 physical strength. Women and other historically disadvantaged groups would want to make sure that rules weren’t biased against their group, and so on. The point is that the outcome of a negotiation would likely be distorted by past injustices, or arbitrary factors, in the preferences of particular individuals. Thus, Rawls seeks a better way to get at justice.

He asks us to imagine individuals who are behind a veil of ignorance getting together to decide on the rules of society. He refers to this as the original position, and structures the original position so that individuals are rational and self-interested but behind a veil of ignorance. The veil of ignorance is such that individuals do not know what their personal characteristics will be. They don't know whether they will be male or female, black or white, high IQ or low IQ, physically strong or weak, musically talented, successful at business, and so on. At the same time, these individuals will be rational and self-interested and know something about human nature and human psychology. In a sense, what Rawls suggests here is that we have to imagine *generic* human beings. They have the abstract features of all human beings in that they are rational and self-interested, and they have general knowledge about how humans behave and interact and how they are affected in various ways, but they have no specific knowledge about who they are or will be when they emerge from behind the veil of ignorance.

According to Rawls, justice is what individuals in the original position would agree to. Justice is what people would choose when 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 two rules would be agreed to in the original position:

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 are "rules of rules" in the sense that they are general principles constraining the formulation of more specific rules. These principles assure that no matter where an individual ends up in the lottery of life (i.e., no matter what kind or degree of intelligence, talents, or physical abilities one has), he or she would have liberty and opportunity. Every individual will have a fair shot at a decent life.

Although Rawls's account of justice has met with criticism, it goes a long way toward providing a framework for envisioning and critiquing just institutions. This discussion of Rawls is extremely abbreviated, as were the accounts of Kant and utilitarianism. As before, we have to stop the dialectic and note that discussion could go off in any number of directions from here. Perhaps the most important thing to keep in mind when claims about rights and justice are made is not to accept them without question. Generally, such claims presume a much more complicated set of concepts and assumptions, and you cannot know whether the claim is worthy until you examine what lies behind it and what its implications are.

Virtue Ethics

One other important tradition in moral philosophy 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 the excellences of human character. A person who has these qualities is one who is capable of functioning 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, whereas the other theories focus 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 is worthy of our admiration?

Virtue theory might be brought into the discussion of computers and information technology and ethics at any number of points. The most obvious is, perhaps, the discussion of professional ethics, where the characteristics of a good computer professional should be considered. 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.

Analogical Reasoning in Computer Ethics

In Chapter 1, we identified one of the goals of computer ethics as understanding the role of computers and IT in constituting situations that pose an ethical dilemma or call for ethical decision making. One very useful way to analyze such situations is to reason by analogy, that is, consider similar (analogous) situations in which there isn't a computer or IT, and then examine whether the absence of the technology makes a moral difference. Sometimes the technology doesn't change the character of the ethical situation; other times it does. Either way, the analogy can be enlightening. Often when we reason by analogy we are able to see things in the analogous case that are relevant to the computer situation but weren't visible because we were focused on the technology. If, on the other hand, the involvement of the technology seems to make a moral difference, then we know there is something about the way in which the technology has constituted the situation that needs to be examined more carefully and linked to a moral concept or theory.

To illustrate, consider a break-in by a computer hacker. This kind of behavior will be discussed more extensively in Chapter 6, but for now let us consider a simple case. A hacker breaks into someone's system, looks around at various files, and copies some of the files that are stored locally. What should we make of this behavior? That is, how should we characterize and evaluate it ethically? Reasoning by analogy, we could consider the similarities and differences between this behavior and that of someone who breaks into an office, then into file cabinets, and then removes paper

files that are of interest. Is there a moral difference between these two kinds of actions? Certainly it is true that the physical movements required to get access to electronic files are quite different from those required to break into an office and into a file cabinet. On the other hand, both sets of actions involve obtaining access to information that an individual had stored with the intention that others would not have access. An interesting difference between the two situations is that in the computer case, the files are still there and available for use by the owner, whereas in the noncomputer case, the files are gone. Does this mean the hacker case is not a case of “theft”? Is it still theft but less harmful or less bad than an ordinary break-in? Or are these morally comparable actions? If so, then “theft” must involve more than depriving the owner of access to what he or she owns. The point is that the analogy helps in teasing out what elements of the case are relevant to a moral assessment and what elements are not. If we cannot find anything morally different about the two cases, then we cannot (with consistency) claim that one type of behavior is morally permissible and the other is not.

Consider a slightly different case with a different analogy. Suppose a hacker is trying to break into systems to see whether he can do it. If he is able to break in, he looks at files but never makes copies. The behavior is mostly about the challenge of breaking in. Is this comparable to walking down a street and testing the doors of every house to see whether they are locked? Suppose someone does this and when they find a door unlocked (a file accessible), they go in and look around. They don’t take anything from the house (file). They simply look at what the owner has put in her or his drawers (what she or he has stored in various files). The question is, is there any difference between these two cases? Is testing to see whether you can get access to computer systems different from testing doors on houses to see whether they are unlocked? From the point of view of the person who is being intruded upon, both types of actions may be felt to be intrusions of privacy and a violation of property rights. Whatever one says about the comparability or noncomparability of these cases, the analogy helps to focus attention on the elements of the action or case that are relevant to a moral evaluation.

Nevertheless, although analogies can be extremely helpful, they have to be used with caution. Reasoning by analogy has dangers that can be avoided only by fully developing the analogy. Analogies are useful because they allow us to draw upon situations or technologies with which we are familiar, situations in which there may be less controversy about right and wrong. This helps us to see rules or principles that might be relevant in the computer situation. The danger is that we may be so taken with the similarities of the cases that we fail to recognize important differences. For example, in arguing about online break-ins and the dissemination of computer viruses, hackers sometimes put forth the argument that they are providing a service by identifying and revealing the flaws and vulnerabilities in computer systems so that they can be fixed. Countering this argument, Eugene Spafford (1992) uses a powerful analogy. He suggests that the hacker’s argument is comparable to arguing that it is morally permissible to set a fire in a shopping mall to show the flaws in the fire protection system. Launching a computer virus on the Internet has some parallels to starting a fire in a shopping mall, but this analogy is so powerful that we might

immediately jump to the conclusion that because one is wrong, the other must also be wrong. We should first ask whether there are any important differences. Some might argue that lighting a fire in a shopping mall puts individual lives at risk, whereas most computer viruses do not. Both actions cause property damage, but the damage done by most computer viruses can be repaired more easily. Thus, when reasoning by analogy, it is important to identify the differences as well as the similarities between the computer and noncomputer cases.

Conclusion

The deep questions and general concerns of ethics that we have discussed in this chapter will continue to come into play in the chapters that follow. The ideas delineated in Chapter 1 as the substance of sociotechnical computer ethics will be brought together with the ethical concepts and theories discussed in this chapter. The goal of Chapter 3 and subsequent chapters will be to analyze ethical issues in IT-configured societies.

As we will see in the next chapters, IT creates a variety of situations that challenge traditional ethical concepts and theories. The dialectic method is enormously helpful in making progress on these challenges. However, it is important to note that the concepts and tools introduced in this chapter are not algorithms for solving moral problems; they are not the be all and end all of practical ethics. They are a starting place. Remember that science is never done either. 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, and why lying and killing are wrong. The stories we tell often get better over time. The stories are retold with new interpretations and in ways that fit the current context. Sometimes accounts get broader (more encompassing) and richer, sometimes more elegant. They are best when they help 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.

Study Questions

1. How do descriptive (empirical) claims and prescriptive (normative) claims differ? Give examples of each kind of claim.
2. Describe a discussion of a moral issue that is currently receiving attention in the media. Identify different claims and arguments that were put forward and defended. List the claims in an order that illustrates a dialectic about this issue, with one claim and argument leading to another claim and another argument, and so on. Are there some claims that are still being presented in the media that have, in your judgment, already been rejected in the dialectic?
3. Explain the difference between “ethics is relative” as a descriptive claim and as a normative claim.
4. What evidence can be used to support “ethics is relative” as a descriptive claim?

5. What are the three problems with “ethics is relative” as a normative claim?
6. What is the basic principle of utilitarianism?
7. What is the difference between an instrumental good and an intrinsic good?
8. Why do utilitarians believe that happiness is the ultimate basis for morality?
9. What is the difference between act-utilitarianism and rule-utilitarianism?
10. What is the major criticism of utilitarianism? Explain it using an example other than the distribution of scarce medical resources.
11. What is the unique characteristic of human beings according to deontologists? How is this quality connected to morality?
12. What is the categorical imperative? Give two examples of violations of the categorical imperative.
13. How can rights be based on deontological theory? How can rights be based on utility theory?
14. What is the veil of ignorance in the original position in Rawls’s social contract theory?
15. What are the two principles of justice in Rawls’s theory?
16. How does virtue ethics theory differ in focus from other theories discussed in this chapter?
17. What is analogical reasoning? Give an example of how it can be used in computer ethics.
18. Why should we always use caution when arguing on the basis of analogies?



Ethics in IT-Configured Societies

CHAPTER OUTLINE

Scenarios

- 3.1 Google in China: “Don’t Be Evil”
- 3.2 Turing Doesn’t Need to Know
- 3.3 Turnitin Dot Com

Introduction: It-Configured Societies

Technology as the Instrumentation of Human Action

Cyborgs, Robots, and Humans

Three Features of It-Configured Activities

Global, Many-to-Many Scope
Distinctive Identity Conditions
Reproducibility

It-Configured Domains of Life

Virtuality, Avatars, and Role-Playing Games
Friendship and Social Networking
Education and Plagiarism Detection

Democracy and the Internet

What Is Democracy?
The Arguments
Is the Internet a Democratic Technology?

Conclusion

Study Questions