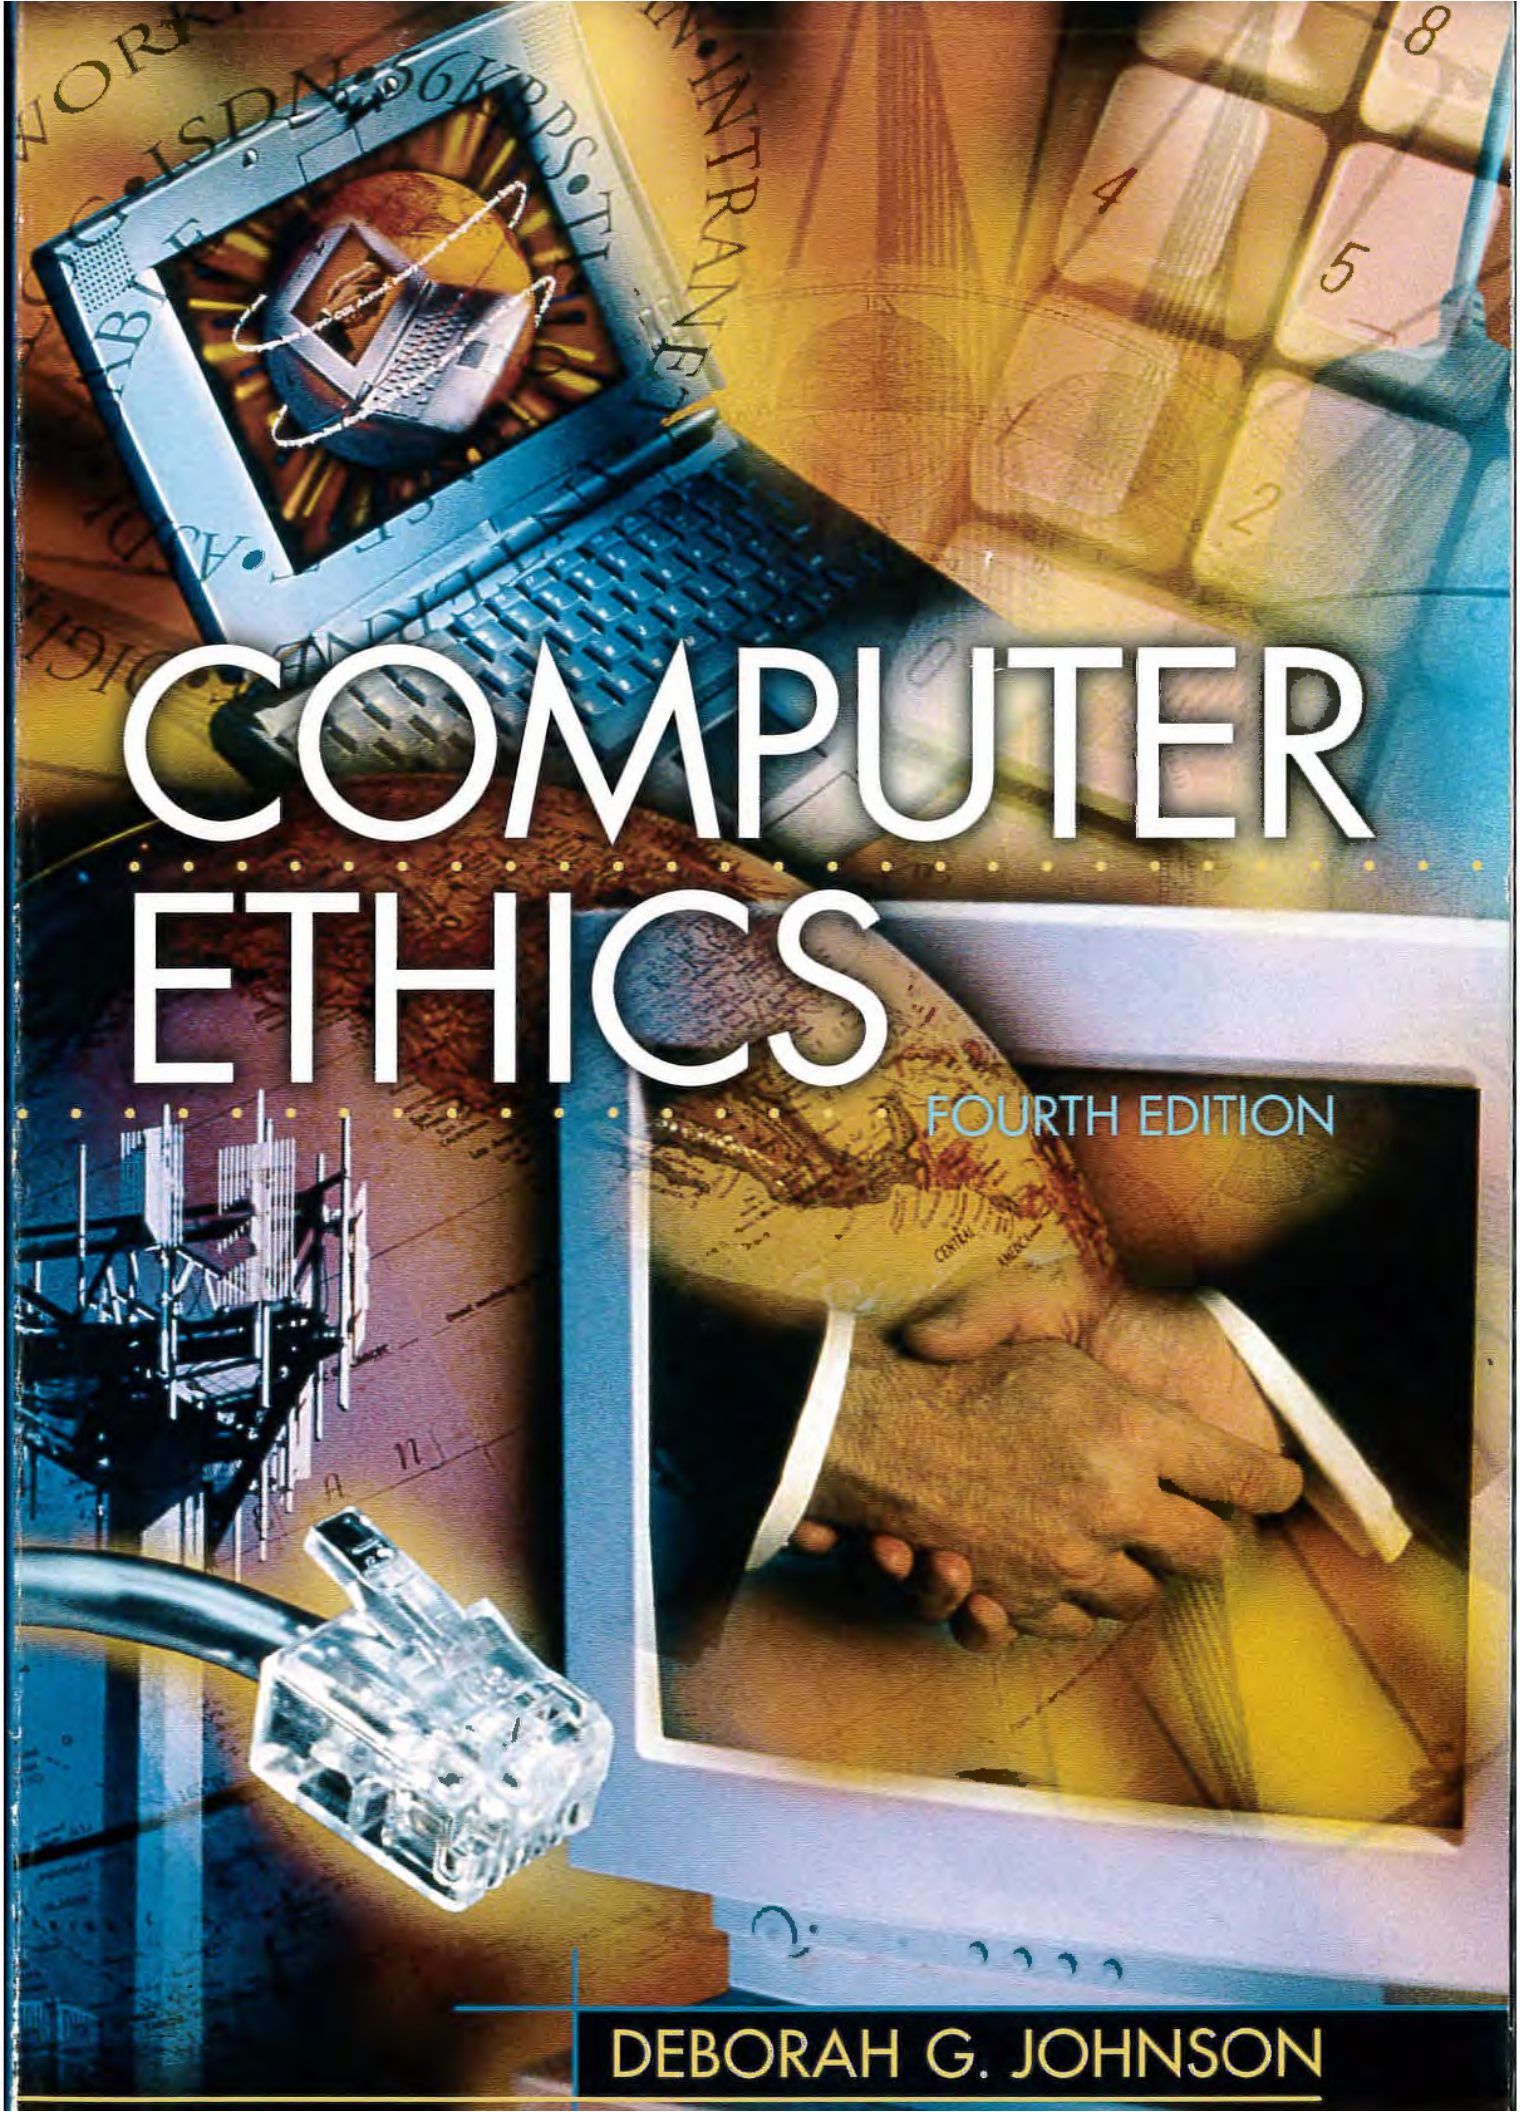

COMPUTER ETHICS

FOURTH EDITION

DEBORAH G. JOHNSON

ACKNOWLEDGMENTS

I am particularly grateful to the University of Virginia's School of Engineering and Applied Sciences for granting me a sesqui leave for the spring semester of 2008. Were it not for this leave, I would not have been able to complete this edition.

During the leave I served as a fellow in residence at two locations of the Erasmus Mundus Master of Applied Ethics program. I spent three weeks in the Department of Philosophy at the Norwegian University of Science and Technology in Trondheim, Norway and three weeks in Linköping University's Centre for Applied Ethics in Linköping, Sweden. During this time I had the opportunity to present drafts of several of the chapters in student seminars, faculty seminars, and public lectures. It is difficult to overstate the value of the honest and helpful responses that I received in both places. And, it was not just the helpfulness of the comments, it was also the warmth and enthusiasm of those whom I met.

I would also like to thank the following reviewers for their helpful suggestions: Karine Barzilai-Nahon, of the University of Washington, Seattle, WA; Xiang Chen, of California Lutheran University, Thousand Oaks, CA; Demetria Enis-Cole, of the University of North Texas, Denton, TX; Peter Madsen, of Carnegie Mellon University, Pittsburgh, PA; and Day W. Radebaugh, of Wichita State University, Wichita, KS.

ABOUT THE AUTHORS

Deborah G. Johnson is the Anne Shirley Carter Olsson Professor of Applied Ethics and Chair of the Department of Science, Technology, and Society at the University of Virginia. Johnson has devoted her career to understanding the connections between ethics and technology. She received the John Barwise prize from the American Philosophical Association in 2004, the Sterling Olmsted Award from the Liberal Education Division of the American Society for Engineering Education in 2001, and the ACM SIGCAS Making A Difference Award in 2000.

Keith W. Miller (Contributor) is a Professor of Computer Science at the University of Illinois at Springfield. His work in software engineering and computer ethics provide complementary perspectives to questions that challenge computer professionals. He is the editor in chief of IEEE Technology and Society, and helped develop the Software Engineering Code of Ethics and Professional Practice. He was named a University of Illinois Scholar in 2000–2003 and received the ACM SIGCAS Outstanding Service Award in 2006.

CHAPTER 1

Introduction to Sociotechnical Computer Ethics

CHAPTER OUTLINE

SCENARIOS

Scenario 1.1 A Virtual Rape

Background:

The following incident took place in the early 1990s and was described by Julian Dibbell in 1993. LambdaMOO is a multiuser designed (MUD) object-oriented program, a complex database maintained inside Xerox Corporation in Palo Alto, California, and open to public access via the Internet. Today there are many more games of this kind with significantly enhanced capabilities. Nevertheless, LambdaMOO remains an intriguing exemplar of the complicated conceptual and ethical issues that arise around computers and information technology.

Case

It happened in the living room in LambdaMOO. The program allows users to create and design the interaction space; a user can create a character with any number of attributes and can build spaces and objects. As users interact with one another as the characters that they have created, they see streams of text, both dialogue and stage descriptions.

One night Bungle entered LambdaMOO and used a subprogram, Voodoo doll, to take control of other characters. Using the Voodoo doll subprogram, Bungle took control of legba and Starspinner, and had the two engage in sadistic actions, with one eating pubic hair and sodomizing the other. Legba and Starspinner were helpless throughout the entire incident. The episode ended when another character, Zippy, used a subprogram to freeze Bungle's commands.

This virtual rape caused enormous ripples across the community of LambdaMOOers. One of the victims, legba, wanted Bungle to be "toaded"—that is, to have his account removed from LambdaMOO. Opinion was divided over what should be done to Bungle. On the evening of the third day after the incident, the users gathered in LambdaMOO to discuss Bungle's fate. There were four arguments: (1) The techno libertarians argued that rape in cyberspace was a technical inevitability, and that a solution would be to use defensive software tools to filter out the offender's words. (2) The legalists argued that Bungle could not legitimately be "toaded" because the MOO had no explicit rules at all; they proposed the establishment of rules and virtual institutions to exercise the control required. (3) The third group believed that only the programmers, or wizards as they are known in MOO, have the power to implement rules. (4) The anarchists, on the other hand, wanted to see the matter resolved without the establishment of social control. There was no agreement between these groups. To Bungle, who joined midway through the conference, the incident was simply a sequence of events in virtual reality that had no consequences for his real life existence.

After weighing the arguments, one of the programmers, the Wizard JoeFeedback, decided to "toad" Bungle and banish him from the MOO. As a result of this incident, the database system was redesigned so that the programmers could make changes based on an action or a petition of the majority of the LambdaMOO community. Eight months and 11 ballots later, widespread participation produced a system of checks and capabilities to guard against the type of violence that had occurred. As for Bungle, he is believed to be reincarnated as the character, Dr Jest.

Did Bungle (or the person controlling Bungle) do anything wrong? Who is responsible for what happened? Should anyone suffer "real-world" consequences?

[Revised from a scenario written for *Computer Ethics 3rd Edition* by Marc Quek Pang, based on J. Dibbell, "A Rape in Cyberspace" *Village Voice* (December 21, 1993), pp. 36–42]

Scenario 1.2 Surprises About Social Networking

Background

Facebook has been wildly popular from its beginning. Although generally identified as a "social networking" site, in recent years users have been surprised by a series of incidents and practices suggesting that the site is much more. A few years ago the company decided to change the architecture of the site so that any time a user added a friend to his or her list of friends, all of the user's friends were alerted to the change. Users didn't like the change and complained so much that Facebook changed the architecture of the site, making the new feature an option but not the default option. A second incident occurred when Facebook introduced a new feature that would generate advertising revenue for the company. The new schema, called Beacon, automated notification of a Facebook member's friends when the member made an online purchase. This advertised the product that the member bought, but it also generated some surprises. One of the stories told in the media was that of a man who was planning to surprise his wife with a ring. The man's plans were ruined when everyone in the man's network was notified of the purchase before the man had a chance to give the ring to his wife. Again users protested and Facebook dismantled Beacon. The third surprise is not associated with any single event. Facebook members have gradually—through a series of incidents—become aware that the site is being used by recruiters and law enforcement agencies to gather information for nonsocial networking purposes. For example, employers search Facebook for information on potential employees, and law enforcement agencies search for information and evidence related to crimes. They look for photos as well as communication related to social events before and after they occur.

Are there any ethical issues here? Did Facebook do anything wrong? Are employers and law enforcement agencies doing anything wrong when they use the site for their purposes?

Hypothetical Situation

Ashley joined Facebook many years ago and now has a site with much information and many pictures from her activities. Shawn, who works for a big company, has also been on Facebook for some years. He was recently employed by a big company in their human resources office. Shawn's job is to interview applicants for jobs; once an applicant has made it through the interview process, Shawn solicits references in writing as well as by phone. Recently Shawn's unit has been brainstorming about better ways to find out about applicants and in a meeting that Shawn didn't attend, the unit decided it would be a good idea to check out applicants on their websites. Shawn is asked to follow up on Ashley who made it through the interview with flying colors; Ashley was highly rated by the interviewers who believe she would be ideal for the job for which she is applying. Shawn easily finds Ashley on Facebook and reports to the interviewers that Ashley appears to party often, and that many of the pictures show her drinking. Fearing that Ashley might not take her job seriously enough, the company decides not to offer the job to Ashley. Ashley is surprised when, weeks later, she discovers that someone else has gotten the job.

Is there anything wrong here?

Scenario 1.3 RFID and Caring for the Elderly

Background

Radio-frequency identification (RFID) is a technology that uses small, passive devices as chips that can be detected from a short distance away from the tag. Some RFID chips are sufficiently small so that the circuitry can be painted directly on an object such as an item of clothing. RFID chips are often used in inventory control. "Computer chips" used to track family pets are RFID chips. A high-profile application of RFID chips is drive-through toll collections and public transportation cards. For almost a decade, a controversial application has been RFID chips placed under people's skin for identification purposes. (Wikipedia, http://en.wikipedia.org/wiki/RFID (accessed January 7, 2007).)

Hypothetical Situation

Kathy Pascal is the legal guardian of her elderly mother, Ada. Ada is in the late stages of Alzheimer's disease, and lives at Golden Oaks, a comfortable nursing home near Kathy's home. Ellen Eiffel, an administrator from Golden Oaks, has contacted Kathy about the possibility of placing an RFID tag under Ada's skin. The tag would be the size of a grain of rice, and Golden Oaks has sensors in many places on their grounds. These sensors record the location of all patients who have an RFID tag whenever they are near a sensor. Ms. Eiffel explains that the RFID tag would help Golden Oaks ensure Ada's safety if and when she started to wander off; it would also help in double checking medical records each time Ada received medicines or therapy. The administrator emphasizes that using the RFID tag would allow Golden Oaks to ensure Ada's safety without confining her to her room. Kathy is sad that her mother requires this kind of marking, but she also sees the advantages as her mother loses more and more of her mental capacity.

What should Kathy do?

INTRODUCTION: WHY COMPUTER ETHICS?

These scenarios illustrate the complex and fascinating character of the ethical and social issues that arise around computers and information technology.[1] Together the scenarios suggest a broad range of issues: Scenario 1.1 presents us with a form of behavior that didn't exist before computers and, thus, requires some analysis just to figure out whether there is any wrongdoing and who did it. Scenario 1.2 raises questions about privacy, uses and abuses of technology, and the obligations of companies to inform their customers about changes in the operation of the business. Scenario 1.3 raises more personal, although no less complicated, issues about how to treat someone who is not capable of making her own decisions, especially when the decision involves a new technology that may affect the kind of care the person will receive. The scenarios suggest that living in a world constituted in part by computers may involve distinctive and especially challenging ethical issues.

The scenarios point to a future that will be powerfully shaped by computers and information technology, assuming that is, that computers and information technology (IT) continue to develop at the speed and with the success it has in the past. If we have any hope of steering the development of future technologies in a direction that is good for humanity, that hope lies in understanding the social and ethical implications of our choices about IT. This book is devoted to just that. The ideas discussed here are intended to provide insight into the social and ethical implications of computers. Those insights should help us think more deeply about the future development of IT.

Although the three scenarios illustrate the range and complexity of ethical issues surrounding IT, some might argue that it is not exactly the technology that poses the ethical challenges but rather the uses of the technology, that is, the humans and human behavior around the technology. In the past, it was common to hear people say that technology is neutral—value neutral—and, therefore, ethics doesn't have anything directly to do with technology. As the old adage goes, "guns don't kill people, people kill people." The field of computer ethics developed when statements of this kind were still quite common and, as a result, much of the literature in the field struggles with questions of the following kind: Why do computers raise ethical issues? What is the connection between ethics and IT? Computer ethicists have struggled with the question of whether IT creates new ethical issues—issues that never existed before—or new versions of old ethical issues, issues that persist over centuries but take on new dimensions as the world changes.

At first glance, it seems that IT creates situations in which common or prevailing moral rules and principles don't seem to apply nor seem helpful in figuring out what one should do. For example, in Scenario 1.1, it takes some analysis just to identify what behavior or whose behavior, if any, could be considered unethical. Because Bungle is a virtual figure, how can it be his behavior? Is it the behavior of the person controlling Bungle? Should we even distinguish the behavior of Bungle from the

[1]Although the focus of this book is broadly on ethics and computers and information technology, because the field of study has traditionally been referred to as "computer ethics," we use "computers" and "computer ethics" in this chapter. In subsequent chapters, we shift to using "information technology" and "IT."

behavior of the person controlling Bungle? Either way, what exactly was the wrongdoing? Was it Bungle's rape of legba and Starspinner? Was it the use of a subprogram to control the behavior of other characters? What moral norms or principles were violated? The prohibition on rape is clear and powerful, but in this case no flesh-and-blood person was raped. Rather, a flesh-and-blood person manipulated a virtual (should we say "fictitious"?) character to enact a text-based rape of another virtual character in front of a community of observers who had no expectation that they would witness such behavior.

The other two scenarios also raise challenging questions. Did Facebook violate the privacy of its members when it introduced changes to the architecture of Facebook? Was this unethical, or simply bad business? Have Facebook members been misled into thinking the site is more private than it is? Has Facebook misled users by offering membership for free when, in fact, Facebook is a for-profit company that must find a way of making money from the site? Are recruiters and law enforcement agencies abusing the site when they use it for other than social networking purposes? As for the nursing home case, although children with elderly parents have often had to make difficult decisions with regard to parents who become incapable of making their own decisions, is the decision about implantation of an RFID chip somehow different than other such decisions? Are such implants dehumanizing and demeaning? Or are the chips the means to a compassionate end?

We will consider these questions in due course but for now, we have to step back and ask a set of larger questions about questions ("meta-questions") regarding the field of computer ethics. Scholars in this field have spent a lot of time trying to understand whether and how ethical issues surrounding IT are distinctive. They have asked whether the issues are so different that new moral theories are needed, or whether traditional theories might be extended to apply. As well, they have considered whether a new kind of methodology is needed for the field. We shall refer to this cluster of issues as the "why computer ethics?" question. The cluster includes: Why does IT create ethical issues? Do we need a special field of study for IT ethics? Why isn't this just applied ethics, plain and simple? In other words, why say that the ethical issues described in Scenarios 1.1–1.3 are computer or IT ethical issues, and not just ethical issues, period? What is the best way to understand and resolve ethical issues that involve IT?

The "why computer ethics?" question is complex. Part of the puzzle has to do with technology in general, because technologies other than computers have also posed complex ethical issues. Consider, for example, all of the concern that was expressed about the power of the atomic bomb during World War II. Should such a powerful tool be created, let alone used? What would it mean for world politics? Or consider, more recently, the public debates about nanotechnology, cloning, stem cell research, and mind-alternating pharmacology. All of these technologies have stirred fear and apprehension as well as fascination and hope. In each case, the literature expressing concern about the new technology has suggested that humanity has acquired a new capacity that takes us into new ethical territory. Part of the "why computer ethics?" question, thus, has to do with technology in general. Why do new technologies give rise to ethical issues? What exactly is the connection between ethics (be it moral theory or moral behavior) and technology?

The other part of the "why computer ethics?" puzzle has to do specifically with IT and whether there is something special about this set of technologies that gives rise to a distinctive kind of ethical issue. On the surface, IT seems to create many more ethical issues than other kinds of technology such as automobiles, electricity, and bridges. Perhaps there is something in particular about IT that disrupts and challenges prevailing moral norms and principles. We will return to this question in a moment.

The "why computer ethics?" question is what we might characterize as a metaquestion, a question about how we are asking our questions. The "why computer ethics?" question calls upon us to step back from engagement with the issues and reflect on our engagement. It asks us to reflect on what we are looking for, and on what we do when we analyze computer ethical issues. On the one hand, this kind of reflection is ideally done after one has some familiarity with the field and some experience analyzing computer ethical issues. For this reason, it would be best to wait until the end of the book to consider the question. On the other hand, an answer to the "why computer ethics?" question also provides a framework for identifying and understanding the issues. As well, an answer to the question points in the direction of an appropriate methodology to use in analyzing computer ethical issues. Hence, we need at least a preliminary answer to the question before we jump into the substance of the topic.

In the next sections, we will provide a preliminary answer to the "why computer ethics?" question, set the scene for subsequent chapters, and suggest a methodology for analyzing computer ethics issues. The answer we will propose recommends that we keep an eye on the connection between ethics and technology in general as the backdrop—the framework—in which computer ethics issues can best be understood.

THE STANDARD ACCOUNT

New Possibilities, a Vacuum of Policies, Conceptual Muddles

A survey of the literature in the field of computer ethics suggests that there is now something like a consensus answer to the "why computer ethics?" question. Computer ethicists seem to accept the general parameters of an account that James Moor provided in a 1985 article entitled, "What is Computer Ethics?" We will refer to this account as the standard account. According to Moor, computers create new possibilities, new opportunities for human action. All three of the scenarios at the beginning of this chapter illustrate this idea. Virtual environments like LambdaMOO didn't and couldn't exist before IT, not, at least, before the Internet had been created. The invention of Facebook created new possibilities for keeping in touch with friends no matter how far away they are or how long ago you last saw each other. Similarly, new possibilities for tracking and monitoring the movements of individuals were created with the invention of RFID. Of course, IT doesn't just create new possibilities for individuals acting alone; new forms of collective and collaborative action are made possible as well. Interest groups on any topic imaginable can form online and take action collectively; companies can operate globally with a relatively high degree of control and speed of action because of the Internet. Families can stay in close communication (maintaining strong bonds) while members are living in geographically distant places.

According to the standard account, these *new possibilities* give rise to ethical questions. Should we pursue the new possibility? How should we pursue it? Who will gain if the possibility is realized? Who will lose? Will pursuit of the new possibility affect fundamental human values? Computer ethicists have risen to the challenge of these new possibilities by taking up tough questions. Is data mining morally acceptable? Should software be proprietary? Are Internet domain names being distributed fairly? Who should be liable for inaccurate or slanderous information that appears in electronic forums? What should we do about child pornography on the Web? Some of these questions have been resolved (or, at least, concern has waned); some have been addressed by law; others continue to be controversial. New questions continue to arise as new possibilities are created. What will Second Life[2] mean? Should we build robots to take care of the elderly as the Japanese are doing? Should we delegate health decisions to artificially intelligent robot doctors? Should we insert intelligence chips in our brains?

That the new possibilities give rise to ethical questions seems to make sense, although we can press further. Why do ethical questions arise from new possibilities? Of course, part of the answer is simply that the new possibilities are "new." But part of the answer is also that new possibilities are not always or necessarily good (or purely good). They can affect different individuals differently. They can be disruptive and threatening to the status quo. The potential for good and ill often comes in a tangled package. Good consequences come along with negative consequences, trade-offs have to be made, and the technology has to be modified in response to political, social, and cultural conditions.

For example, virtual reality systems have enormous potential for good. Aside from the rich, entertainment value of gaming, virtual systems used for scientific modeling and simulation help in understanding the world and in training. But virtual systems could also lead, some fear, to a world in which individuals escape into fantasy worlds and have difficulty dealing with the "real world" of flesh and blood people. Similarly, a world in which RFID is used to monitor and track those who are hospitalized could mean a world in which the elderly are much better cared for than they are now, or it could mean the elderly have less and less human contact and nurses and doctors become deskilled and lose first-hand knowledge of illness and aging.

Thus, according to the standard account of computer ethics, the field's raison de trios (reason for being) is to evaluate the new possibilities from an ethical perspective. To be sure, the implications of adoption and use of a particular technology can and should be examined from a variety of perspectives, including economics and politics, but the ethical perspective is especially important because it is normative. When it comes to economics and politics, the point is often to describe and predict the likely consequences of adopting a new technology. This informs but does not address whether the new technology *should* be adopted. Ethical analysis considers the should-question and how a new possibility fits (or doesn't fit) moral values, notions, and practices.

[2]Created in 2003, Second Life is a popular 3-D virtual world site in which users interact through avatars. Because of the advanced capabilities of the site, users sometimes strongly identify with their avatars and become intensely involved in their virtual lives.

Moor (1985) describes the task of computer ethics as that of filling policy vacuums. According to Moor, when computers create new possibilities, there is a vacuum of policies. The new possibilities take us into uncharted territory, situations in which it is unclear what is at issue or which moral norms are relevant. Moor's notion of a policy vacuum captures the uncertainty that often surrounds the invention and adoption of new technologies. Here an example from the early days of computer technology illustrates Moor's point. When the first computers were installed, individuals began storing files on them, but there were no institutional or legal policies with regard to access and use. From our perspective today, it may seem obvious that most computer files should be treated as personal or private property, but the status of computer files was initially unclear (in part because the first computers were large mainframes located in buildings and owned by companies, agencies, and universities). Thus, when remote access became possible and hackers began roaming around and trying to get access, the moral and legal status of the files on mainframe computers was unclear. Whether or not hackers were committing crimes was unclear. Were they stealing? Perhaps, but the files that hackers accessed (and copied) were not removed. Were they trespassing? Hackers who gained access were nowhere near the physical location where the files were stored. As already indicated, at the time there were no laws explicitly addressing access to computer files. In Moor's terms, there was a policy vacuum with regard to the status of acts involving access to computer files. A new possibility had been created and there was a policy vacuum.

On Moor's account, the task of computer ethics is to fill policy vacuums, and he acknowledges that the task is far from easy. Filling the policy vacuum involves sorting out what Moor refers to as conceptual muddles. To illustrate a conceptual muddle, consider another case from the early days of computing, computer software. When computer software was first created, the challenge was to figure out how best to conceptualize it. The problem had to do with fitting computer software to prevailing intellectual property law; copyright and patent seemed the best possibilities. Copyright law specifies that abstract ideas cannot be copyrighted, only expressions of ideas. Typically this means expressing an idea in a written language. Patent law also prohibits ownership of abstract ideas, as well as laws of nature, and mathematical algorithms. Because abstract ideas are the building blocks of science and technology, giving an individual ownership has the potential to significantly dampen progress in the technological arts and sciences. New inventors would have to get permission from a private owner to use one of the building blocks. When it came to software it wasn't clear whether a copyright on a computer program would be granting ownership of an expression of an idea or the building blocks of the electronic world. In patent law the issue was even trickier because patent law specifies that abstract ideas, laws of nature, and mental steps cannot be owned. Although enormously large and complex, software can be thought of as a series of mental steps. That is, in principle a person can go through the steps in a program and mentally do what the program specifies. If someone were granted ownership of mental steps, then they could legally prohibit others from going through those steps in their minds. This would interfere with freedom of thought.

The question of whether to grant copyright or patents for computer programs was, then, deeply linked to the conceptualization of computer programs. That is, the policy vacuum couldn't be filled without a conceptualization of software. Could software be characterized as an expression of ideas? an application of abstract ideas? Could it be understood as something other than mental steps or mathematical algorithms? Or did a whole new set of laws have to be created specifically for computer software? If so, what should the new laws look like? Again, the conceptual muddle had to be sorted out in order to fill the policy vacuum.

In summary, then, according to the standard account of computer ethics: (1) ethical issues arise around IT because IT creates new possibilities for human action and there is a vacuum of policies with regard to the new possibilities, (2) the task of computer ethics is to evaluate the new possibilities and fill the policy vacuums, and (3) a significant component of this task is addressing conceptual muddles.

An Update to the Standard Account

The standard account has been extremely useful in moving the field of computer ethics forward for the last two decades. Nevertheless, over these years, a number of factors have changed. IT has changed and so have computer ethicists, at least in the sense that they have acquired a good deal of experience in analyzing IT ethical issues. At the same time, a new field of study has developed, science and technology studies (STS). This new field has provided insights into the relationship between technology and society, insights that are relevant to understanding how ethical notions and practices shape, and are shaped by, technology. These factors suggest that it is time for an update to the standard account.

To begin the update, notice that much of what has been said about IT ethics seems to apply quite readily to ethical issues involving other new technologies. Other new technologies also create new possibilities for human action, and the new possibilities lead to ethical questions about whether and how to pursue the possibilities. Should I donate my organs for transplantation? Should employers be allowed to use urine or blood tests to determine whether employees are using drugs? Should our food supply be genetically modified? Each of these questions arose when a new technology was developed, and the new possibility created an option for human action that hadn't existed before. Because of features of the new technology, prevailing moral norms or rules either didn't apply, or didn't apply neatly to the new possibility. For example, having a supervisor watch employees as they worked on an assembly line of a manufacturing plant was a standard part of such work, but when urine and blood tests for illegal drugs were developed and adopted by employers, it wasn't clear whether this was an extension of acceptable workplace watching practices or an inappropriately intrusive step into the private lives of individuals. Although there was a huge body of law relating to employer and employee rights, the applicability of the law to urine and blood testing was unclear. Is it comparable to watching employees at home? Is it like asking about an employee's race, sexual preference, or political beliefs? Is drug testing comparable to watching an employee work? So, there was a policy vacuum and a conceptual muddle. The point is that the standard account can be used to explain ethical issues arising around new technologies in general, and is not specific to IT ethics.

Moor's account is, then, an account of "new technology ethics"; something more is needed to make it an account of computer ethics. Of course, its broad applicability is not a reason to reject the account. It seems to be an accurate account of new technology ethics; we just have to keep in mind that it is not specific to IT.

Another, perhaps more subtle problem with the standard account is that the emphasis on "newness" may skew the kind of analysis that is done. The focus of attention is on one, and only one, stage in the lifecycle of technology, the stage in which it is first introduced. This directs attention away from, and largely blinds us to, other stages, especially the ongoing role of IT in constituting the social and moral world. IT is an ongoing part of the world we live in. Indeed, it is challenging to identify a domain of life in which IT doesn't play a role.

The focus on newness suggests that computer ethics issues arise when the technology is first introduced; the issues get resolved when the policy vacuums are filled and the conceptual muddles sorted out, and that is that. The reality is quite different. For one thing, policy vacuums sometimes go unfilled or they get filled, but in ways that perpetuate continuous struggle or tension over the policy. Sometimes policy vacuums are resolved with bad policies, policies with negative or undesirable consequences. In any of these cases, ethical analysis can have an important role in critiquing policies that have already formed, pointing to their misfit with notions of justice or responsibility or good consequences. Moreover, even when a policy issue gets resolved and gets resolved well, because IT constitutes the social and moral world, it is still important to draw attention to the role of IT in shaping moral practices.

The emphasis on newness in the standard account leads to other related problems. Because IT is no longer new, many who take up IT ethical issues (indeed, many readers of this book) will not have experienced a world without computers. Yes, novel applications, tools, and systems continue to be developed, but they are developed in a context in which people are already familiar with the technology. The technology already has meaning and users already have well-developed expectations. In other words, people already have conceptual models of the technology and how it works; they have knowledge that informs how they approach and use new applications. Also, there are already policies regulating the use of computer technology, policies that are extended to new applications and systems when they are introduced. Hence, it no longer seems appropriate to frame computer ethics as a field focused exclusively on the newness or novelty of IT.

Yet another reason for shifting away from the focus on newness is to avoid a presumption that seems to accompany it. When we focus on IT when it is new, we tend to think of the technology as arriving intact and being plopped into society where it is taken up and has an impact. This suggests that the technology came out of nowhere, or that it was developed in isolation from society and then "introduced." Many believe that technology is developed in protected environments such as laboratories, garages, and universities, as if the inventors or designers were following out some logic of nature. But this is an incomplete account; technologies are always developed in a social world. Laboratories, universities, and even garages are embedded in an existing culture, complete with systems of support, real-world constraints, and socially shaped ideas. What a product or tool looks like—the features it includes, what it makes possible—has everything to do with the social context *in* which it was created and the context *for* which it was created.

IT systems are designed to do certain tasks, to fit into particular environments, or fit the needs or desires of particular users. The process by which they are designed—who is involved, who has a say, who is funding the project—powerfully shapes the particular features a technology comes to have, and who it serves or doesn't serve. The invention and design context is filled with legal requirements, economic incentives, cultural attitudes, and consumers with particular profiles. Moreover, after a technology is first created, it is often modified either by users who find "work-arounds," or by developers who see that their product is being rejected, or by others who see something in it of value but see an alternative way to develop the general idea. Post-it notes were born when a failed glue found a new use. More often than not, successful technologies have gone through a long period of development with many missteps and unexpected turns along the way.

Inventors live and work in particular places and at particular periods in history, and this has significant effects on the range of possibilities open to them. The garage where Steve Jobs and Steve Wozniak assembled the first Apple computers was near Hewlett-Packard in Palo Alto where Wozniak had worked and Jobs had attended some lectures. The garage had electricity, and Wozniak had been trained as an electronic engineer. The new computers they designed used the existing technologies of a mouse and on-screen icons. When the two made their first batch of computers, a nearby computer hobby store bought them. Apple computers were products not only of the unique talents and skills of Jobs and Wozniak; they were also the products of the circumstances and possibilities that were available at that time and place.

The focus on newness can, then, blind us to the fact that the technologies we have today are not the only possible, or the best possible, technologies. Under different conditions and with different design and innovation contexts, we could have different technologies. When the possibility of different sorts of technology is pushed out of sight, IT ethicists miss an important opportunity for ethics to play a role in the design of IT systems. [Note: Several IT ethicists have not been blinded in this way, and have seized the opportunity to both focus on issues in the design of IT systems and be involved themselves in the design. Helen Nissenbaum, for example, helped to design a program called TrackMeNot that helps protect users' privacy when they use Google to search the Web. She and others have developed an approach to IT ethics that is referred to as *value sensitive design*.]

So, the standard account does not provide an account of distinctively "computer" or "IT" ethical issues; it gives an account of how new technologies, in general, involve ethical issues. Because IT is relatively new, and new applications continue to evolve, IT falls under the account. In putting the emphasis on newness, the standard account tends to push out of sight other stages in the lifecycle of technologies. Before adoption and use is the design stage, and here computer ethics can play a role in identifying ethical issues in both design processes and design features. After design and introduction, IT continues to contribute to the configuration of social arrangements, social practices, and social institutions. IT is part of, and shapes, many domains of life including government, education, politics, business, identity, and relationships. As well, the lifecycle of IT includes manufacturing, marketing, distribution, and disposal. The lens of ethics should be brought to bear on all of these stages in the lifecycle of IT.

The central focus of the rest of this book will be on the role of IT in constituting the social and moral world. For this purpose, it will be helpful to adopt what we will refer to as the sociotechnical systems perspective.

THE SOCIOTECHNICAL SYSTEMS PERSPECTIVE

In the last thirty years, a rich literature focused on the relationships among science, technology, and society has developed. The literature is part of a new field of study with undergraduate majors and graduate programs, several journals, and professional societies. The field is called "STS," referring either to "science and technology studies" or "science, technology, and society." We will develop a foundation for IT ethics using STS insights, concepts, and theories. For this, a brief review of the major claims of STS is essential.

STS literature is diverse, complex, and richly textured, so the description to follow is necessarily a simplification. To provide a quick overview of the core ideas in STS, we can think of STS as identifying three mistakes that should be avoided in thinking about technology. Parallel to each of the three mistakes is a recommendation as to how we should think about technology and society.

Reject Technological Determinism/Think Coshaping

STS cautions against adoption of a view referred to as "technological determinism." Although multiple definitions and forms of technological determinism have been articulated, technological determinism fundamentally consists of two claims: (1) technology develops independently from society, and (2) when a technology is taken up and used in a society, it determines the character of that society.

The first claim of technological determinism usually involves thinking that technological development follows scientific discoveries or follows a logic of its own, with one invention building on the next. Technological determinists may even think that technological development has a kind of natural evolution with each development building on previous developments. This view of how technology develops goes hand-in-hand with the belief, mentioned earlier, that inventors and engineers work in isolation. They work, it is supposed, in laboratories in which all that matters is manipulating the processes and materials of nature. Technological development is understood to be an independent activity, separate from social forces.

STS scholars reject this claim. They argue that scientific and technological development is far from isolated and does not follow a predetermined or "natural" order of development. The character and direction of technological development are influenced by a wide range of social factors including: the decisions a government agency makes to fund certain kinds of research; social incidents such as a war or terrorist attack that spark interest and effort to produce particular kinds of devices (e.g., for security); market forces that feed development in some areas and bring others to a halt; the legal environment, which may constrain innovation in certain areas and facilitate it in others; and cultural sensibilities that lead to objectionable meanings associated with certain technologies and desirable meanings for others.

Consider, for example, the enormous investment that the U.S. government (through its National Science Foundation) has made in the development of nanotechnology. The NSF receives far more requests for funding than it can grant, so making investments in nanotechnology means that the NSF will not make investments elsewhere. Consider how regulatory standards for automobile safety and standards for fuel efficiency have influenced the design of automobiles. And consider the debates over stem cell research and the requirement that researchers obtain the informed consent of subjects on whom they experiment. These are all elements that have shaped the technologies that are currently in use.

To be sure, nature has to be taken into account in technological development. Nature cannot be made to do just anything that humans want it to do. Nevertheless, nature does not entirely determine the technologies we get. Social factors steer engineers in certain directions and influence the design of technological devices and systems. Thus, the first tenet of technological determinism—that technology develops in isolation and according to its own logic—should be rejected outright.

According to the second tenet of technological determinism, when technologies are adopted by societies or particular social groups, the adoption brings about—determines—social arrangements and patterns of social behavior. In other words, when a society adopts a particular technology, it adopts a form of life, patterns of behavior. Perhaps the most famous statement of this was historian Lynn White's claim (1962) that from the invention of the stirrup came feudal society. He was suggesting that the adoption of the stirrup changed the nature of warfare and slowly but surely led to a society in which serfs were dominated by aristocrats. In IT ethics, a parallel type of claim is made about the Internet and democracy. Certain writers have suggested that when countries adopt the Internet, it is just a matter of time before democracy will reign; once, that is, individuals in any society have access to the Internet and all the information it makes available, those individuals will want democracy and democratic social institutions. This is an expression of technological determinism in the sense that it implies that a technology will determine the political structure of a country.

Although STS scholars reject outright the first claim of technological determinism, their response to the second claim is more complicated. The problem is that when we say that technology determines society, we are forgetting that the technology has been socially shaped; social factors and forces have influenced the development and design of the technology. As already discussed, STS studies show that the technologies we have today are products of highly complex and contingent social processes. Thus, the problem with claiming that technology determines society is that "determines" is too strong a term. Social factors affect the design, use, and meaning of a technology, and in this respect society can push back and reconfigure a technology, making it into something its designers never intended. Consider here how Facebook users pushed back and pressured the company to change the architecture back to what it was. The point is that although technology shapes society, it does not determine it.

Technology develops through a back-and-forth process that involves what is technologically possible and how society responds to the possibilities, pursuing some possibilities, rejecting others, and not even noticing others. So, technological determinism is not wrong insofar as it recognizes technology as a powerful force in shaping society; it is wrong to characterize this as "determining" society. Society and technology shape each other.

In effect, the STS counter to each tenet of technological determinism is the same; society influences technology. Recognition of the societal influences on the development of technology leads to an outright rejection of the first claim of technological determinism (that technology is developed in isolation) and a modification to, or weakening of, the second tenet (that technology determines society). The positive recommendation emerging out of this critique of technological determinism is that we acknowledge that technology and society cocreate (coshape; coconstitute) one another. The mantra of STS scholars is that technology shapes and is shaped by society, that society shapes and is shaped by technology.

In the previous section, in critiquing the standard account, we mentioned that the account seemed to frame computer technology as developed, and then "introduced" as if it came in some sort of predetermined form and was simply discovered. Framing the development of technology in this way commits the mistake of technological determinism. It suggests that users have only one choice: either reject or accept the technology as delivered. Nothing could be farther from the truth; IT is developed to fit into particular environments; users are often able to shape the technology by customizing settings, demanding changes from developers, and choosing between alternative products, and so on. Users also shape computer technology through the meaning they associate with it and through the behavior with which they engage the technology.

Reject Technology as Material Objects/Think Sociotechnical Systems

The second major insight that STS theory provides involves the rejection of another presumption that people often make about technology. They think and speak as if "technology" refers to physical objects or artifacts. According to STS scholars, this is at best misleading, and at worst constitutes a false conception of technology. To be sure, artifacts (human-made material objects) are components of technology, but artifacts have no meaning or significance or even usefulness unless they are embedded in social practices and social activities. This can be seen in a number of different ways. First, technologies do not come into being out of nowhere; they are created by intentional human activity and, as already described, shaped by social forces. This is true whether we think about a simple artifact created by a single individual fashioning natural materials (say, a person carving a stone into an arrowhead), or we think about an extremely complex artifact such as a mass-produced computer that requires elaborate social organization. Producing a computer involves the organization of people and things into manufacturing plants, mining of materials, assembly lines, distribution systems, as well as the invention of computer languages, education and training of individuals with a variety of expertise, and more. In other words, technology is a *social product.*

However, technology is not just the outcome of social activity, it is also socially constituted (it *is* social). The artifactual component of technology (the physical object) can function only as part of a social system. The mere existence of an artifact doesn't do anything. Consider, for example, a workplace monitoring system adopted by a corporation for use by those who supervise employees working on computers. The system is a social product in the sense that it took many people, organized in various ways and working intentionally, to develop the system. However, the system doesn't work once the code is written. The company making the system will have had to figure out how to legally "own" the system (via patent or copyright or trade secrecy) before they make it available. Documentation will have to be written. The system will have to be advertised or marketed, and it will have to be distributed. If customers buy the system, users have to be trained; users have to learn how to adapt the system to their particular needs (kind of work being done, number of employees, kind of output desired); and users have to learn how to interpret and use the data produced by the system. The customer may have to write a new policy regarding the system; they may have to inform workers and obtain their consent. In short, in order for the workplace monitoring system to work, the software has to be embedded in a set of social practices. The thing we call a "workplace monitoring system" consists not just of software but the combination of software and human arrangements and social practices; these all work together to make a functional system.

It is misleading, then, to think of technology as merely artifacts or of IT systems as merely software. STS theorists recommend that we think of technology as sociotechnical systems (Hughes, 1994). A frequent series of TV commercials for a cell phone company in the United States features a horde of technicians, operators, and other personnel who follow around customers of that company. That "horde" is a lighthearted illustration of exactly this STS concept: A cell phone is not just the artifact that you put to your ear, talk into, and listen for. A cell phone is the combination of the artifact and a network of people arranged in various ways to produce a complex of results.

Recognition that technology is not just artifacts, but rather artifacts embedded in social practices and infused with social meaning, is essential to understanding the connection between ethics and IT. Traditionally, ethics has been understood to be almost exclusively about human behavior and human action. Ethicists have not traditionally focused on technology, perhaps, because they believed that technology was simply material objects dropped into the world ready-made. Because material objects were thought simply to be products of nature, they were seen as neutral, and there seemed to be no point to ethical reflection about them. This is precisely the danger of thinking about technology as material objects. It pushes out of sight the fact that people and artifacts are intertwined, that people are influenced by artifacts, and that artifacts are shaped by humans. For ethicists to fail to see the role of technology in morality is to fail to see a powerful force shaping the moral questions confronting human beings. For engineers, inventors, and computer experts not to see the social practices that constitute technological systems they develop is to be blind to the significance and implications of what they are doing.

So, STS scholars reject the idea that technology is material objects, and entreat us always to think of technology as sociotechnical systems (combinations of things and people). As already indicated, this doesn't mean that artifacts are unimportant; they are enormously important. The material world powerfully shapes what people can and cannot do. However, we will be misled if we look only at artifacts. In fact, it could be argued that it is impossible to understand a technology by looking at the artifact alone. This would be like trying to understand the chess piece called "the rook" without knowing anything about the game of chess (the rules of the game, the goal, or other chess pieces). Yes, you can describe the shape and dimensions and the material of which the chess piece is made, but you cannot fully understand what a rook "is" without reference to the game of chess. It is the same for a workplace monitoring device, a word processor, or a data-mining tool: You cannot understand what they are merely by focusing on the code.

Reject Technology as Neutral/Think Technology Infused with Values

The third mistake identified in the STS literature is to think that technology is value neutral. Perhaps the most influential work on this topic is Langdon Winner's 1986 piece, "Do artifacts have politics?" Winner draws attention to the relationship between technology and systems of power and authority, arguing that particular technologies cannot exist or function without particular kinds of social arrangements. He argues that adoption of a particular technology means adoption of a particular social order. His example is that of nuclear power: Nuclear power necessitates a complex, hierarchical system of decision making; the production and distribution of nuclear power is achieved by social arrangements in which decisions are coordinated and someone is in charge. Experts of various kinds make decisions at various nodes in the organization. Contrast this with windmills that operate with a decentralized form of authority; each individual who has a windmill can decide how to operate the windmill and what to do with the power that is produced. Similarly, transportation by train requires a centralized system of organization, whereas bicycling is decentralized.

In explaining this relationship between technologies and patterns of authority and decision making (which may seem quite deterministic), Winner provides a powerful example of how an artifact can enforce social biases and privilege individual agendas. He describes how Robert Moses intentionally designed the bridges of Long Island, New York (built in the 1930s) to be at a height that would not allow public buses to go under the underpasses. This constrained bus routes and prevented poor people (largely African Americans) living in the city from getting to the beaches. In the 1930s poor people didn't have cars, so the only way they could reach the beaches during the heat of the summer was by public transportation. This account of Moses's intentions has been challenged, but whether or not it was consciously intended by Moses, the account illustrates that the height of bridges can constrain access to certain areas and thus can reinforce a race and class system.

The story is intriguing because it illustrates how a material object can be value-laden. One is tempted to say that social hierarchy was embedded in the materiality of the bridges. Of course, it isn't the physical structures alone that produced the social arrangement. It was the combination of the bridge's size and height, the size of other physical objects (buses of a particular size, the location of the beaches in relation to the city), and a set of social practices including the practice of going to public beaches, thinking of people in racial categories, and much more. The combination constituted a race-biased arrangement. Still, all of these parts constitute the soicotechnical system of which the physical bridge is a part, and that system was infused with social and moral values.

Winner can be interpreted as slipping into the mistake of technological determinism. He seems to be suggesting that a technology—the bridges of Long Island—determined the social order. Hence, it is important to remember here that the problem with technological determinism is not that it is wrong about technology "shaping" or "influencing" social arrangements; technology does shape and influence social behavior. Technological determinism goes too far in claiming that the technology determines the social arrangements. Here we see that the social arrangement was produced by the combination of the height of the bridges, the size of buses, preexisting social arrangements, and ideas about race and social hierarchy; a change in any one of these elements might have changed the result.

SOCIOTECHNICAL COMPUTER ETHICS

The three STS recommendations provide the foundation for what we will call "sociotechnical computer ethics." The payoff of using this approach will become clearer as we move from issue to issue and chapter to chapter, but we can demonstrate some of its value here if we return to the scenarios at the beginning of this chapter. The "Virtual Rape" case will be taken up in Chapter 3 but a closer look at the Facebook and RFID cases will get us started. Our analysis will be limited because we have not yet explored ethical concepts and theories. We will use Facebook to illustrate the STS recommendations and the RFID case to demonstrate how the sociotechnical perspective helps in ethical decision making.

The story of Facebook's development goes right to the heart of the first STS theme in the sense that Facebook was not the "next logical development in the natural evolution of IT"; Facebook didn't come out of nowhere. It was created by Mark Zuckerberg while he was at Harvard and thought it would be fun to create something that would support social interactions among students. Whether he was conscious of his knowledge or not, Zuckerberg used his understanding of patterns of interaction among college students; he designed a system that would fit into that world; he intentionally designed a system that would enhance and extend prevailing patterns of interaction. As the system began to be used, it affected social relations by facilitating students in finding out about one another and interacting more frequently via the system. Among other things, Facebook allows individuals to communicate asynchronously, with different people than they might otherwise and, of course, independent of where they are located. In these respects, Facebook shapes the nature of friendship. However, it would be an overstatement to say that Facebook "determines" social relationships. Facebook shapes, and is shaped by, the nature of relationships.

Perhaps the second STS lesson—not to think of technology as material objects—doesn't even need emphasizing to Facebook users because they think of the site not just as a material object or piece of software, but as a "social" networking site. They are aware that what makes Facebook work is not just lines of code, but users putting up content, browsing, and communicating with one another. The surprises described in Scenario 1.2 reinforced this idea because they made users painfully aware of the human actors involved in making the system work. Systems operators and administrators had made decisions to change the architecture of the system, and later they decided to change the architecture back to its original form. So, users were confronted with the fact that the system is not simply lines of code; it is partly lines of code, but the lines of code are written and maintained by programmers who take direction from administrators who respond to a variety of stakeholders, including users. Facebook is a sociotechnical system with many human and nonhuman components.

As a social networking site, Facebook is far from neutral. It is designed to facilitate social networking. Once again, the surprises to the users illustrate this point. Users didn't want a system that would send out information to their friends every time they made a change in their list of friends. Although the system makes individuals quite transparent to their friends, the Beacon schema bumped up against many users' desire from some sort of privacy about shopping. These incidents show that users' values and preferences were in tension with Facebook's values. The Facebook company wants a system that makes money; users want a system that makes some, but not other, information available to their friends. Changes in the architecture change the values embedded in the system.

Facebook illustrates the three STS themes and recommendations. Still, you might ask, how do these STS recommendations help us when it comes to ethical issues? That is, how does the sociotechnical systems perspective help in the analysis of IT ethical issues? The short answer is that the perspective gives us a fuller, more accurate, and richer understanding of situations in which moral questions and dilemmas arise. We can illustrate this by focusing on Scenario 1.3 in which an individual must make a decision about whether to have an RFID device implanted in her mother.

The first step is to keep in mind that RFID is a sociotechnical system, not simply a material object. Those who developed the chips to be implanted in patients saw a real-world context and a set of practices that might be improved by the use of the chips. So they designed the chips for use in the hospital context. In the design process, the developers had to pay considerable attention to how things are done in hospitals—who does what when; they had to take into account the interests of various stakeholders including, and especially, the hospital and patients. The developers had to think about what might go wrong and what their liability would be if something did go wrong. They would have had to think through how the device could be inserted, by whom, and under what conditions; how data on patients could be displayed; where the displays would be located; who would monitor the displays; who would need to be trained; and so on. All of this is to say that the RFID device at issue is a social product and a sociotechnical system. It is created by people with interests and targeted for other people organized in a particular setting. The system is a combination of material chips together with social practices involving implantation of the tag, display of the data produced by the tag, interpretation of the data, and responses to the data.

When installed at a facility, the RFID chip system becomes a component shaping the hospital environment. Patients may move about more freely because the device will inform staff as to their medical condition and whereabouts, staff will be trained to read and interpret displays, staff may be assigned more patients, and the physical architecture of hospitals may change because patients can move about more. Perhaps the most important lesson in this for Kathy Pascal in Scenario 1.3 is that her decision doesn't involve just insertion of a chip into her mother; it involves her mother being enmeshed in a complex sociotechnical system with many components.

Although Kathy should ask about the materials used in the chip and whether there are risks to implantation of those materials, she should also ask about how the chip will be implanted, how data will be received, and how decisions will be made using the data. She will want to compare the treatment her mother is likely to get with or without the chip system. Although her mother may be freer to move about, does this mean she will have less contact with people during her day? Is the science behind the detection devices (that will monitor her mother's medications) reliable? And so on.

Yes, the sociotechnical systems perspective seems to generate more questions than someone without the perspective would have thought to ask. Although this may seem a burden, it is unavoidable that better decisions involve taking into account more factors. Yet the sociotechnical system perspective doesn't just expand the range of factors to be taken into account; it helps in identifying or articulating particular kinds of concerns, and reveals new opportunities for resolution or intervention. For example, suppose Kathy is already concerned about the chip being demeaning and disrespectful of whatever autonomy her mother has. To figure out whether the chip will have this effect or not, if Kathy focuses on the chip alone, she will get nowhere. On the other hand, once she recognizes the chip as part of a larger system, she is led to gather information about the whole system and this may help her evaluate whether the system is demeaning or not. It depends on how her mother is treated during the surgical implantation, how the data is used by hospital staff, whether implantation means less human interaction with hospital personnel, and so on.

It may be that Kathy cannot do anything about the composition of the system; that is, her decision may be a matter of simply saying "yes" or "no" to the implant. But that yes/no decision can be made more wisely after the sociotechnical systems perspective reveals a range of options for hospital administrators and the systems developers. For example, if they find the device is being rejected because patients (or their loved ones) find it demeaning, they may be able to identify different nodes in the system where changes might be made. It may not be the chip itself that has to be changed or abandoned but rather a change in the implantation procedure, in the user interface, or in the training of hospital staff. Changes in one of these nodes will change the nature of the system and may alter perceptions or attitudes toward the system.

In summary, the sociotechnical systems perspective provides a richer account of situations in which ethical decisions are made, one that may help in articulating moral concerns as well as revealing additional avenues for addressing ethical questions and issues.

MICRO- AND MACRO-LEVEL ANALYSIS

One final distinction will set the scene for the next chapter. In ethics, a distinction is often made between macro- and micro-level analysis. Micro-level analysis focuses on individuals, their choices, and their behavior. In the scenario just discussed, Kathy Pascal is faced with a decision, a personal decision, and in this respect the scenario raises a micro-level question. What should Kathy do? What are her responsibilities? What factors should she take into account? By contrast, macro issues are generally focused on groups or organizations or even countries, and they are generally concerned with policies, rules, or systems. What should the hospital's policies and procedures be with respect to RFID devices? What privacy policies should the United States or the European Union adopt? Should employers monitor employee e-mail? Should software be proprietary?

This distinction is important as we identify, articulate, and answer ethical questions. However, the relationship between the two levels of analysis is complex. Issues at one level of analysis impact issues at another level. For example, in the Facebook scenario, we described Facebook at the macro level (that is, we described the Facebook company changing its policies) and then described a hypothetical situation that posed a micro ethical issue: What should Shawn do with information he finds on Facebook? Sometimes micro-level questions are answered by referring to a rule established at the macro level. For example, if we focus on an individual who breaks into a computer system and gains unauthorized access and ask the micro-level question whether the individual did anything wrong, we may answer that question simply by referring to a macro-level rule or law. The following sentence is an example of the interaction of micro and macro ethical analysis: "The hacker was wrong to gain unauthorized access because it is illegal."

Because the sociotechnical perspective frames technology as a system, it seems to draw more attention to macro-level issues. However, as we saw in our analysis of Kathy's situation, macro analysis enhances micro-level analysis. Thus, the sociotechnical systems perspective is compatible with, and useful to, both levels of analysis.

RETURN TO THE "WHY COMPUTER ETHICS?" QUESTION

We can now re-ask the "why computer ethics?" question: Why is a book (or course or field of study) focused on computer ethics needed? As was noted at the start, there are two questions here: Why a book on technology and ethics? Why a book specifically on computers or IT and ethics? Both questions can now be answered.

Technology is a part of human activity. It makes a difference in the way we live and act, it shapes the moral issues we confront and the courses of action (the options) that are available to us, and it affects the decisions we make, individually and collectively. The better we understand technology and how it shapes and is shaped by morality, the better our choices and decisions are likely to be. That is the answer to the first question. The second question arises because all technologies are not the same. Different technologies affect human activity and forms of life differently. The field of computer ethics focuses specifically on the role of IT in constituting the moral world. General studies of technology and ethics inform IT ethics, and IT ethics informs the broader study of technology and ethics. So the two work together.

Yet another question at the beginning of the chapter can now be addressed. What, we asked, is the relationship between "ethics" and "IT ethics"? Why isn't IT ethics just ethics? The sociotechnical systems perspective reveals that all social activities and practices are, in part at least, shaped by technology, so whether ethicists have recognized it or not, technology has, effectively, always played a role in moral practices and moral thought. Moral philosophy is focused on human action and social arrangements, and technology has always been intertwined with both. In this respect IT ethics is part of ethics, but in IT ethics we highlight and pay special attention to the role of IT as one of many elements that come into play in moral practices, decisions, and outcomes. Thus, it seems best to say that IT ethics is a subfield of ethics.

This particular subfield of ethics happens to be a lively scholarly area at the moment. Philosophers, computer scientists, sociologists, lawyers, and others are debating many issues surrounding IT and IT ethics. This book cannot adequately discuss all of the issues currently of interest in the literature. For example, a recent and intensely debated theory called "Information Ethics" insists that all information objects, including humans, should be afforded ethical respect because all information should be protected from entropy. If embraced, this theory would have broad-ranging implications for IT ethics. However, in this book, we will focus on more established ethical theories. [For those who want to follow the "information ethics" stream of analysis, L. Floridi's "Information ethics: On the philosophical foundation of computer ethics" *Ethics and Information Technology* 1: 37–56, 1999, is a good starting place.]

Conclusion

According to the standard account, ethical issues arise around IT because IT creates new possibilities for human action, and there is a vacuum of policies with regard to the new possibilities. The task of computer ethics is, then, to evaluate the new possibilities and fill the policy vacuums. A significant component of this task is addressing conceptual muddles. The standard account has been shown here not to be wrong but insufficient because it does not provide an account of distinctively "IT" ethical issues. It provides an account of how new technologies, in general, create ethical issues, and because IT is relatively new and new applications continue to evolve, IT falls under the account. The emphasis on newness was shown to be problematic for other reasons, as well. In particular, it puts the focus on IT when it is new and first introduced, and, thus, skews attention away from the ongoing role of the technology in structuring our moral conditions. Moreover, the standard account can blind us to the importance of the design of computer systems and, hence, to opportunities to change the moral conditions in various environments by changing the technology.

Drawing on insights from the field of STS, we have proposed an alternative account of IT ethics that we refer to as "sociotechnical IT ethics." Although sociotechnical IT ethics goes a long way toward supplementing the standard account and avoiding its pitfalls, the sociotechnical approach does not, by any means, make the

task of IT ethics easy. The sociotechnical perspective emphasizes that the social and technological elements are interwoven, and it tells us that we are misleading ourselves if we think we can entirely disentangle these elements. This should make our conclusions more realistic and better informed, but it will also require more nuanced and sophisticated analysis.

Sociotechnical IT ethics has never, to our knowledge, been explicitly attempted on the scale of a book. In the chapters that follow we show that IT applications are sociotechnical systems, that is, combinations of software, hardware, and social practices, and that these combinations help to constitute the world in which human beings—individually and collectively—act. Viewing these systems as sociotechnical systems provides the foundation for richer analyses and more options for addressing ethical issues in IT.

Study Questions

1. What is encompassed in the question "why computer ethics?" In other words, what more specific questions are included in the broad question?
2. Give a concise summary of Moor's standard account of computer ethics.
3. What is a policy vacuum? Give an example, and explain it using the standard account.
4. What is a conceptual muddle, and how do they get in the way of filling policy vacuums? Illustrate with an example.
5. What is wrong with the standard account of computer ethics? Identify at least two criticisms, and explain.
6. What are the two tenets of the view referred to as "technological determinism"? What is the STS critique of each tenet?
7. What is wrong with thinking of technology as merely material objects?
8. What is a sociotechnical system?
9. Choose a familiar technology and describe its values.
10. What is the difference between micro-level ethical issues and macro-level ethical issues? Give an example of each.
11. Why is the study of ethics and technology needed? Why is the study of ethics, and IT in particular, needed?