CHAPTER 1

The Internet and Ethical Values

The end [of ethics] is action, not knowledge.

—Aristotle[1]

Light: © sakkmesterke/Shutterstock; Gavel: © Maksim Kabakou/Shutterstock

Many decades have passed since the first communications were transmitted over a fledgling global network, which would later be called the *Internet*. At the time, few would have predicted the Internet's explosive growth and persistent encroachment on our personal and professional lives. This radically decentralized network has been described in lofty terms as empowering and democratizing. It has lived up to this ideal by creating opportunity for many new voices with extraordinary reach. Although the claim that the Internet will revolutionize communications may be hyperbole, there is no doubt that the Internet has the potential to magnify the power of the individual and fortify democratic processes.

Many governments, however, are clearly threatened by some of this decentralized power and they have sought to impose some centralized controls on this anarchic network. The United States has attempted to regulate speech through the ill-fated Communications Decency Act and to restrict the use of encryption technology through its key recovery scheme. More draconian regulations have been imposed by countries like Iran, China, and Saudi Arabia. The Net and its stakeholders have steadfastly resisted the imposition of such controls, and this has led to many of the tensions and controversies we consider throughout this text.

Although the control of technology through law and regulation has often been a futile effort, "correcting" technology with other

technology has been more effective. The regime of law has had a hard time suppressing the dissemination of pornography on the Internet, but blocking software systems that filter out indecent material have been much more successful. This reflects the Net's paradoxical nature—it empowers individuals and allows them to exercise their rights such as free speech more vigorously, but it also makes possible effective technical controls that can undermine those rights.

Although the primary axis of discussion in this text is the ethical issues that surface on the Internet, we must devote attention to these related matters of cyber governance and public policy. Thus, we explore in some detail the tensions between the radical empowerment that the Net allows and the impulse to tame this technology through laws and other mechanisms.

Because this is a text about ethics, about *acting* well in this new realm of cyberspace, we begin by reviewing some basic concepts that will enrich our moral assessment of these issues. Hence, in this introductory chapter our purpose is to provide a concise overview of the traditional ethical frameworks that can guide our analysis of the moral dilemmas and social problems that arise in cyberspace.

More important, we also elaborate here on the two underlying assumptions of this work: (1) the *directive* and architectonic role of moral ideals and principles in determining responsible behavior in cyberspace and (2) the capacity of free and responsible human beings to exercise some control over the forces of technology (technological realism). Let us begin with the initial premise concerning the proper role of cyberethics.

Cyberethics and the "Law of the Horse"

An ethical norm such as the imperative to be truthful is just one example of a constraint on our behavior. In the real world, there are other constraints, including the laws of civil society or even the social pressures of the communities in which we live and work. There are many forces at work limiting our behavior, but where does ethics fit in?

This same question can be posed about cyberspace, and to help us reflect on this question we turn to the framework of Larry Lessig. In his highly influential book, *Code and Other Laws of Cyberspace*, Lessig first describes the four constraints that regulate our behavior in real space: law, norms, the market, and code.

Laws, according to Lessig, are rules imposed by the government that are enforced through *ex post* sanctions. There is, for example, the

complicated IRS tax code, a set of laws that dictates how much taxes we owe the federal government. If we break these laws, we can be subjected to fines or other penalties levied by the government. Thanks to law's coercive pedagogy, those who get caught violating tax laws are usually quick to reform.

Social norms, on the other hand, are expressions of the community. Most communities have a well-defined sense of normalcy, which is reflected in their norms or standards of behavior. Cigar smokers are not usually welcome at most community functions. There may be no laws against cigar smoking in a particular setting, but those who try to smoke cigars will most likely be stigmatized and ostracized by others. When we deviate from these norms, we are behaving in a way that is socially "abnormal."

The third regulative force is the market. The market regulates through the price it sets for goods and services or for labor. Unlike norms and laws, market forces are not an expression of a community and they are imposed immediately (not in *ex post* fashion). Unless you hand over $2 at the local Starbucks, you cannot walk away with a cup of their coffee.

The final modality of regulation is known as architecture. The world consists of many physical constraints on our behavior; some of these are natural (such as the Rocky Mountains), whereas others are human constructs (such as buildings and bridges). A room without windows imposes certain constraints because no one can see outside. Once again "enforcement" is not *ex post*, but at the same time, the constraint is imposed. Moreover, this architectural constraint is "self-enforcing"—it does not require the intermediation of an agent who makes an arrest or who chastises a member of the community. According to Lessig, "the constraints of architecture are self-executing in a way that the constraints of law, norms, and the market are not."[2]

In cyberspace we are subject to the same four constraints. Laws, such as those that provide copyright and patent protection, regulate behavior by proscribing certain activities and by imposing *ex post* sanctions for violators. It may be commonplace to download and upload copyrighted digital music, but this activity breaks the law. There is a lively debate about whether cyberspace requires a unique set of laws or whether the laws that apply to real space will apply here as well, with some adjustments and fine tuning. Judge Frank Easterbrook has said that just as there is no need for a "law of the horse," there is no need for a "law of cyberspace."[3]

Markets regulate behavior in various ways—advertisers gravitate to more popular websites, which enables those sites to enhance services; the pricing policies of the Internet service providers determine access to

the Internet; and so forth. It should be noted that the constraints of the market are often different in cyberspace than they are in real space. For instance, pornography is much easier and less expensive to distribute in cyberspace than in real space, and this increases its available supply.

The counterpart of architectural constraint in the physical world is software "code," that is, programs and protocols that make up the Internet. They, too, constrain and control our activities. These programs are often referred to as the "architectures of cyberspace." Code, for example, limits access to certain websites by demanding a username and password. Cookie technology enables e-commerce but compromises the consumer's privacy. Sophisticated software is deployed to filter out unsolicited commercial email (or spam). In the long run, code may be more effective than law in containing spam, which rankles many users.

Finally, there are norms that regulate cyberspace behavior, including Internet etiquette and social customs. For example, spamming and hacking were always considered "bad form" on the Internet, and those who did it were chastised by other members of the Internet community. Just as in real space, cyberspace communities rely on shame and social stigma to enforce cultural norms.

But what role does ethics play in this neat regulatory framework? Lessig apparently includes ethical standards in the broad category he calls "norms," but in our view cultural norms should be segregated from ethical ideals and principles. Cultural norms are nothing more than variable social action guides, completely relative and dependent on a given social or cultural environment. Their validity depends to some extent on custom, prevalent attitudes, public opinion, and myriad other factors. Just as customs differ from country to country, the social customs of cyberspace could be quite different from the customs found in real space. Also, these customs will likely undergo some transformation over time as the Internet continues to evolve.

The fundamental principles of ethics, however, are metanorms; they have universal validity. They remain the same whether we are doing business in Venezuela or interacting in cyberspace. Like cultural norms, they are prescriptive; but unlike these norms, they have lasting and durable value because they transcend space and time. Ethics is about (or should be about) intrinsic human goods and the moral choices that realize those goods. Hence, the continuity of ethical principles despite the diversity of cultures.

Our assumption that ethics and customs (or cultural norms) must be kept distinct defies the popular notion of ethical relativism, which often equates the two. A full refutation of that viewpoint is beyond the scope of our discussion here. But consider the reflections of the

contemporary philosopher, Phillippa Foot, about cultural diversity. She carefully argues that while it is obviously wrong to assume the exact identity between people of different cultures; there is certainly a great deal that all human persons share in common with one another. The human person is intrinsically relational. Therefore, we all need love and affection, the cooperation of others, and an opportunity to live in community. Human beings simply cannot flourish without these things. When there is isolation and constant divisiveness or an absence of friendship and loving kindness, human fulfillment is damaged or impeded. According to Foot, we are not referring to arbitrary standards

> if we think of some moral systems as good moral systems and others as bad. Communities as well as individuals can live wisely or unwisely, and this is largely the result of their values and the codes of behavior that they teach. Looking at these societies, and critically also at our own, we surely have some idea of how things [will] work out

based on values.[4]

None of this by any means invalidates Lessig's framework. His chief insight is that "code and market and norms and law together regulate in cyberspace as architecture and market and norms and law regulate in real space."[5] Also, according to Lessig, "Laws affect the pace of technological change, but the structures of software can do even more to curtail freedom. In the long run the shackles built by programmers could well constrain us more."[6] This notion that private code can be a more potent constraining force than public law has significant implications. The use of code as a surrogate for law may mean that certain public goods or moral values once protected by law will now be ignored or compromised by those who develop or utilize this code. Moreover, there is a danger that government itself will regulate the architectures of cyberspace to make it more controllable. We have already seen this happen in countries such as Iran and China. In the hands of the private or public sector, the architectures of cyberspace can have extraordinary regulatory power.

Thus, Lessig's model is quite instructive and we rely on it extensively in the pages to come. However, I would argue that the model would be more useful for our purposes if greater attention were given to the role of fixed ethical values as a constraining force. But how do these values fit with the other regulatory forces?

Before we can answer this question we must say something about the nature of those values. The notion that there are transcendent moral values grounded in our common human nature has a deep tradition in the history of philosophy. It is intuitively obvious that there are basic human goods

that contribute to human well-being or human flourishing. Although there are several different versions of what these goods might be, they do not necessarily contradict each other. Some versions of the human good are "thin," whereas others are "thick." James Moor's list of core human goods includes life, happiness, and autonomy. According to Moor, *happiness* is "pleasure and the absence of pain," and *autonomy* includes those goods that we need to complete our projects (ability, security, knowledge, freedom, opportunity, reason). Individuals may rank these values differently, but all human beings attribute value to these goods or "they would not survive very long."[7]

Oxford philosopher John Finnis offers a thicker version of the human good. He argues persuasively for the following list of intrinsic goods: life, knowledge, play (and skillful work), aesthetic experience, sociability, religion, and practical reasonableness (which includes autonomy). According to Finnis, participation in these goods allows us to achieve genuine human flourishing. They are opportunities for realizing our full potential as human beings, for being all that we can be. Hence, the master principle of morality: one's choices should always be open to *integral human fulfillment*, the fulfillment of all persons and communities. None of our projects or objectives provides sufficient reason for setting aside or ignoring that responsibility.

For both Moor and Finnis, then, the ulitmate source of moral normativity is these intelligible, authentically *human* goods, which adequately explain the reasons for our choices and actions, and overcome the presumption of subjectivism. Morality can begin to claim objectivity because this collection of basic human goods is not subjective, that is, subject to cultural differences or individual whims.

The ultimate good, the human flourishing of ourselves and of others, should function as a prescriptive guidepost of enduring value, serving as a basis for crafting laws, developing social institutions, or regulating the Internet. Because this moral ideal is rather lofty, its application to policy making can be difficult. As a result, we are also guided by intermediate ethical principles, such as the Golden Rule; do to others what you would have them do to you. Similarly, one could be guided by Kant's second version of the categorical imperative: "Act so that you treat humanity always as an end and never merely as a means."[8] From these principles one can derive more *specific core moral values* about murder, theft, or lying. These principles can function as more practical guidelines for moral decision making and enable us to pursue the basic human goods in a way that respects our fellow humanity. According to Finnis, our fundamental responsibility is to respect each of these human goods "in each person whose well-being we choose to affect."[9]

We contend, therefore, that these intelligible goods, intrinsic to human persons and essential for human flourishing, along with basic moral principles (such as the Golden Rule) that protect those goods should play an architectonic or *directive role* in the regulation of cyberspace. They should guide and direct the ways in which code, laws, the market, and social norms exercise their regulatory power. The value of human flourishing is the ultimate constraint on our behavior in real space and in cyberspace. Accordingly, we have enhanced Lessig's model as depicted in **FIGURE 1-1**.

To illustrate our point about the role of these supreme ethical values and how they can be translated into the actual world of our experience, let us consider the regulatory impact of code. There are responsible and irresponsible ways of developing code that constrain behavior. Blocking software systems has become a common way of protecting young children from pornography, as will be discussed in Chapter 3. Those who write this code have developed proprietary blocking criteria, and as a rule they do not reveal these criteria or the specific sites that are blocked. In some cases, sex education or health-related sites are filtered out along with the pornography. If this is done inadvertently, the software should be fixed; if it is done deliberately, parents should be informed that the scope of the filtering mechanism is broader than just pornography. One could certainly make the case that parents should know what the blocking criteria are in order to make an informed judgement about the suitability of this software. Failure to reveal this information is tantamount to disrespecting parental autonomy. As a result,

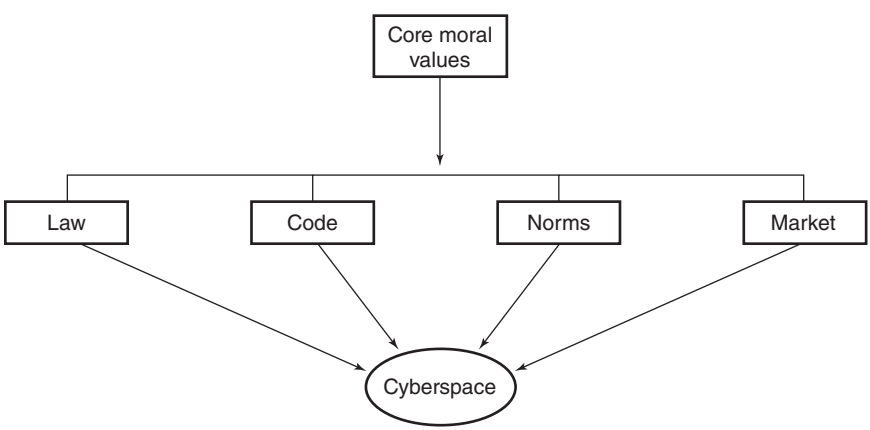

FIGURE 1-1 Constraints on Cyberspace Activities (adapted from Professor Lessig's framework).

one could argue that when the criteria are obscured for some ulterior agenda, the code is not being deployed in a responsible manner that is consistent with the core good of autonomy.

I am not suggesting that this is a clear-cut matter or that moral principles can provide all the answers to proper cyberspace regulations. And I am not making a judgment about whether law or code is the more effective constraint for cyberporn. I am simply claiming that those who write these programs or formulate laws to regulate cyberspace should rely on ethics as a guide. Code writers must be responsible and prudent enough to incorporate into the new architectures of cyberspace structures that preserve basic moral values such as autonomy and privacy. Further, government regulations of cyberspace must not yield to the temptation to impose excessive controls. Regulators, too, must be guided by high moral standards and respect for basic human values such as freedom and privacy. The code itself is a powerful sovereign force, and unless it is developed and regulated appropriately, it will surely threaten the preservation of those values.

The role of morality should now be quite evident: it must be the ultimate regulator of cyberspace that sets the boundaries for activities and policies. It should direct and harmonize the forces of law, code, the market, and social norms so that interactions and dealings there will be measured, fair, and just.

Iron Cage or Gateway to Utopia?

Although most of us agree that some constraints will need to be imposed on the technologies of networking and computing that have come to pervade the home and workplace, there is legitimate skepticism about anyone's ability to control the ultimate evolution and effects of these technologies. Are our attempts to regulate cyberspace merely a chimera? Are we too trammeled by the forces of technology, or are we still capable of exercising sovereignty over the code that constitutes the inner workings of the Internet?

Some philosophers as we observed in the Preface have long regarded technology as a dark and oppressive force that menaces our individuality and authenticity. These technology determinists see technology as an independent and dehumanizing force beyond humanity's capacity to control it. The French philosopher Jacques Ellul presents a disturbing vision of technology in his seminal work, *The Technological Society*. His central argument is that *technique* has become a dominant and untranscendable human value. He defines technique as "*the totality*

of methods rationally arrived at and having absolute efficiency (for a given stage of development) in *every* field of human activity."[10] According to Ellul, technique is beyond our control; it has become autonomous and "fashioned an omnivorous world which obeys its own laws and which has renounced all tradition."[11] For Ellul, modern technology has irreversibly shaped the way we live, work, and interact in this world.

Ellul was not alone in advancing such a pessimistic outlook on technology. Max Weber coined the term *iron cage* to connote how technology locks us in to certain ways of being or patterns of behavior. And Martin Heidegger saw technology not merely as a tool that we can manipulate but as a way of "being in the world" that deeply affects how we relate to that world. But is it really so that technology forces us into this "iron cage" and into a more fragmented, narrow-minded society dominated by a crude instrumental rationality?

In contrast to the bleak outlook of Ellul and Heidegger, we find technology neutralists who argue that technology is a neutral force, completely dependent on human aims and objectives. According to this viewpoint, technologies are free of bias and do not promote one type of behavior over another. Technology *is* only a tool, and it does not compromise our human freedom or determine our destiny in any appreciable way; it is up to us whether this powerful force is used for good or ill purposes.

Some go even further and embrace a sort of "technological utopianism" that regards certain technologies as making possible an ideal world with improved lifestyles and workplaces. This optimistic philosophy assumes that humanity can eradicate many of technology's adverse effects and manipulate this tool effectively to improve the human condition.

The philosophy of technological neutralism (or, for that matter, utopianism) seems problematic for several reasons. Technology does condition our choices with certain "givens" that are virtually impossible to fully overcome. Langdon Winner describes this as a process of reverse adaptation or "the adjustment of human ends to match the character of the available means."[12]

However, in our view, it is also an exaggeration to claim that computer and network technology locks us into a virtual but inescapable iron cage. The middle ground between these extreme positions is *technological realism*, which holds that "although technology has a force of its own, it is not independent of political and social forces."[13] Technological realism acknowledges that technology has reconfigured our political and social reality and that it does influence human behavior in particular ways. To some extent, this notion is echoed in Lessig's work. He argues that we fail to see sometimes how code is an instrument of social and political control. Code is not neutral. Most often,

embedded within code are certain value decisions that define the set of options for policy problems.

Nonetheless, although technology determines to some degree how we live and work, we still have the capacity to redirect or subdue it when necessary. In effect, we can still shape and dictate how certain technological innovations will be deployed and restrained, particularly when there is a conflict with the common good or core human goods. Our human freedom is undoubtedly attenuated by technology's might and its atomizing tendencies, but it is not completely effaced. We can still choose to implement systems and develop code in ways that protect fundamental human rights such as autonomy or privacy. We can be liberated from the thralldom of privacy-invading code by developing new code that enhances privacy.

Beyond any doubt, technology and its counterpart—instrumental rationality—are dominant forces in this society that exert enormous pressures on us to make choices and behave in certain ways. But as Charles Taylor points out, one can find throughout history pockets of concerted opposition to oppressive technologies. Further, the chances for such successful resistance are greatly enhanced when there is some common understanding about a particular threat or imperilment, such as the threat to our ecology that occupied us during the 1970s. Perhaps the same common consciousness will emerge about the threat to personal privacy, and this will provide yet another impetus for human choice to trump the dominating forces of information technology. Although we should not be overly optimistic about our freedom and our capacity for resisting infatuation with new technology, we must recognize that we still have *some* degree of freedom in this world. Thus, we agree with Taylor's assessment: "We are not, indeed, locked in. But there is a slope, an incline in things that is all too easy to slide down."[14]

How then do we avoid this fatal slide? This brings us to our next topic of discussion—the importance of cultivating and sustaining a moral point of view as one deliberates about how to constrain behavior on the Internet through market forces, code, norms, or law.

Ethical Values and the Digital Frontier

We avoid this slide and its accompanying perils only if we conscientiously adopt the moral point of view as we evaluate technological capabilities and make decisions about the ground rules of the digital frontier. How can we characterize this moral point of view? According

to Kenneth Goodpaster, it can be seen "as a mental and emotional standpoint from which all persons have a special dignity or worth, from which the Golden Rule derives its worth, and from which words like *ought* and *duty* derive their meaning."[15] This is quite consistent with our earlier claim that the fundamental moral imperative is the promotion of human flourishing, both in ourselves and in others.

Several distinct types of ethical reasoning have been associated with the moral point of view, and they provide us with the basic principles that serve as a moral yardstick or "compass" that can assist us in making normative judgements. Our discussion here is concise; for the interested reader it can certainly be amplified by many other books on ethical theory or on applied ethics.[16] We consider several models of ethical reasoning based on moral frameworks emphasizing the maximization of social utility, natural rights, contract rights, and moral duties.

The fact that there are several different theories embodying the moral point of view does not contradict our assumption regarding the core human goods that form the basis of a unifying moral framework. All of these theories recognize such goods in one form or another. Kant embraces the principle that we must respect humanity in all our choices and actions, although he might define *humanity* differently from Finnis. And rights-based theories discuss core human goods in terms of protection of human rights such as the rights to life, liberty, and the pursuit of happiness. The utilitarian approach emphasizes happiness, and although it may have a hard time standing on its own, it can be complemented by other theories to form a more comprehensive framework.

All of these theories are worth our careful consideration. Each represents a valuable perspective from which complex moral issues can be assessed and reflected upon. They help us to engage in the critical moral analysis necessitated by the thorny dilemmas that are beginning to surface all over the Internet.

Before we discuss these theories, it is worth pointing out that modern ethical frameworks fall under two broad categories: teleological or deontological. *Teleological* derives from the Greek *telos,* which means *goal* or *end*. These theories argue that the rightness or wrongness of an action depends on whether it brings about the end in question (such as happiness). *Deontological* theories, on the other hand, consider actions to be intrinsically right or wrong—their rightness or wrongness does not depend in any way on the consequences that they effect. These frameworks emphasize duty and obligation (*deon* is the Greek word for *duty*).

Utilitarianism

Utilitarianism is a teleological theory, and it is by far the most popular version of consequentialism. Classic utilitarianism was developed by two British philosophers, Jeremy Bentham (1748–1832) and John Stuart Mill (1806–1873). According to this theory, the right course of action is to promote the general good. This general good can also be described in terms of "utility," and this principle of utility is the foundation of morality and the ultimate criterion of right and wrong. *Utility* refers to the net benefits (or good) created by an action. According to Frankena, utilitarianism is the view that "the sole ultimate standard of right, wrong and obligation is the *principle of utility* or *beneficence,* which says quite strictly that the moral end to be sought in all that we do is *the greatest possible balance of good over evil* (or the least possible balance of evil over good)."[17] Thus, an action or policy is right if it produces the greatest net benefits or the lowest net costs (assuming that all of the alternatives impose some net cost).

It should be emphasized that utilitarianism is quite different from ethical egoism. An action is right not if it produces utility for the person performing that action but for all parties affected by the action. With this in mind we might reformulate the moral principle of utilitarianism as follows: persons ought to act in a way that promotes the maximum net expectable utility, that is, the greatest net benefits or the lowest net costs, for the broadest community affected by their actions.

On a practical level, utilitarianism requires us to make moral decisions by means of a rational, objective cost/benefit analysis. In most ethical dilemmas there are several possible alternatives or courses of action. Once one has sorted out the most viable and sensible alternatives, each one is evaluated in terms of its costs and benefits (both direct and indirect). Based on this analysis, one chooses the alternative that produces the greatest net expectable utility, that is, the one with the greatest net benefits (or the lowest net costs) for the widest community affected by that alternative.

A concrete example illustrates how cost/benefit analysis might work. Let us assume that a corporation has to make a policy decision about random inspection of employee email. This might be done as a routine part of a performance review as a means of checking to make sure that workers are using email only for work-related purposes and are not involved in any untoward activities. This practice is perfectly legal, but some managers wonder if it's really the right thing to do; it seems to violate the privacy rights of employees. Rightness in the

utilitarian ethical model is determined by consequences that become transparent in a cost–benefit analysis. In this case, the managers might face three options: email messages are not inspected on a routine basis and are kept confidential (unless some sort of malfeasance or criminal activity is suspected); email messages are inspected regularly by managers, but employees are informed of this policy and reminded of it every time they log in to the email system, so that there is no expectation of privacy; or email is regularly but surreptitiously perused by managers with employees uninformed of the company policy. Which of these alternatives promotes the general good, that is, produces the greatest net expectable utility?

TABLE 1-1 provides an idea of how this analysis might work out. It becomes clear from this exercise that it is difficult to objectively calculate the diffuse consequences of our actions or policies and to weigh them appropriately. And herein lies a major obstacle in using

TABLE 1-1 Illustrative Cost/Benefit Analysis

	Costs	Benefits
1. Keep email confidential	Lack of control over employees; difficult to prevent misuses of email; email could be used for various personal reasons without company knowledge.	Maintains morale and an environment of trust and respect for workers; protects personal privacy rights.
2. Inspect email with employees informed of policy	Violates privacy rights; diminishes trust and impairs morale; workers less likely to use email if communications are not confidential—instead they will rely on less efficient modes of communication.	Prevents misuse along with inappropriate comments about superiors and fellow workers via email; workers know the risks of using email; they are less likely to use email for personal purposes.
3. Inspect email surreptitiously	Same as option 2, but even more loss of trust and morale if company policy is uncovered.	Better chance to catch employees doing something wrong such as transmitting trade secrets; perfectly legal.

this approach. Nonetheless, there is value in performing this type of analysis; it induces us to consider the broad consequences of our actions and to take into account the human as well as the economic costs of implementing various technologies.

Although this theory does have certain strengths, it is also seriously flawed in some ways. Depending on the context, utilitarianism could be used to justify the infliction of pain on a small number of individuals for the sake of the happiness or benefits of the majority. There are no intrinsically unjust or immoral acts for the utilitarian, and this poses a problem. What happens when human rights conflict with utility? Can those rights be suppressed on occasion for the general good? There is nothing in utilitarianism to prevent this from happening, as long as a cogent and objective case is made that the benefits of doing so exceed the costs. The primary problem then is that this theory lacks the proper sensitivity to the vital ideals of justice and human rights.

Contract Rights (Contractarianism)

Another mode of reasoning that exemplifies the moral point of view is rights-based analysis, which is sometimes called *contractarianism.* Unlike utilitarianism, contractarianism is a deontological theory. It looks at moral issues from the viewpoint of the human rights that may be at stake. A *right* is an entitlement or a claim to something. For instance, thanks to the Fourth Amendment, American citizens are entitled to protection from unwarranted search and seizure in the privacy of their homes. In contrast to the utilitarian view, the consequences of an action are morally irrelevant for those who support contractarianism. Rights are unequivocally enjoyed by all citizens, and the rights of the minority cannot be suspended or abolished even if that abolition will maximize social welfare.

An important distinction needs to be made between positive and negative rights. Possession of a *negative right* implies that one is free from external interference in one's affairs. Examples of negative rights include the right to free speech, the right to property, and the right to privacy. Because all citizens have a right to privacy in their homes, the state cannot interfere in their affairs by tapping their phone calls unless it has demonstrated a strong probability that laws are being broken.

A *positive right,* on the other hand, implies a requirement that the holder of this right be provided with whatever one needs to pursue one's legitimate interests. The rights to medical care and

education are examples of positive rights. In the United States, the right to health insurance funded by the government may still be a matter of debate, but the right to education is unequivocal. Therefore the state has a duty to educate children through the twelfth grade. If everyone had a "right" to Internet access, there would be a correlative duty on the part of the government to provide that access for those who could not afford it.

Rights can be philosophically grounded in several ways. Some traditional philosophers such as Locke and Rousseau and the contemporary social philosopher John Rawls claim that we have basic rights by virtue of an implicit social contract between the individual and civil society. Individuals agree to a contract outside of the organized civil society that stipulates the fundamental principles of their association including their rights and duties. Rights are one side of a quid pro quo—we are guaranteed certain rights (e.g., life, liberty, and the pursuit of happiness) as long as we obey the laws and regulations of civil society. This contract is not real but hypothetical. According to Kelbley, "we are not discussing facts but an ideal which rational individuals can embrace as a standard to measure the moral nature of social institutions and efforts at reform."[18]

According to this perspective, moral reasoning should be governed by respect for these individual rights and by a philosophy of fairness. As Ken Goodpaster observes, "Fairness is explained as a condition that prevails when all individuals are accorded equal respect as participants in social arrangements."[19] In short, then, this rights-based approach to ethics focuses on the need to respect an individual's legal, moral, and contractual rights as the basis of justice and fairness.

The problem with most rights-based theories is that they do not provide adequate criteria for resolving practical disputes when rights are in conflict. For example, those who send spam (unsolicited commercial email) over the Internet claim that they are exercising their right to free speech, but many recipients argue that spam is intrusive, maybe even a form of trespass. Hence, they claim that the transmission of spam is an invasion of their property rights. The real difficulty is how we adjudicate this conflict and determine which right takes priority. Rights-based theories are not always helpful in making this determination.

Moral Duty (Pluralism)

The next framework for consideration is not based on rights, but on duty. The moral philosophy of Immanuel Kant (1724–1804), which can be found in his short but difficult masterpiece on ethics,

Fundamental Principles of the Metaphysics of Morals, is representative of this approach. It assumes that the moral point of view is best expressed by discerning and carrying out one's moral duty. This duty-based, deontological ethical framework is sometimes referred to as *pluralism.*

Kant believed that consequences of an action are morally irrelevant: "An action performed from duty does not have its moral worth in the purpose which is to be achieved through it but in the maxim by which it is determined."[20] According to Kant, actions only have moral worth when they are done for the sake of duty. But what is our duty and how is it derived? In Kant's systematic philosophy our moral duty is simple: to follow the moral law which, like the laws of science or physics, must be rational. Also, as is the case for all rational laws, the moral law must be universal, because universality represents the common character of rationality and law. And this universal moral law is expressed as the categorical imperative: "I should never act except in such a way that I can also will that my maxim should become a universal law."[21] The imperative is "categorical" because it does not allow for any exceptions.

A *maxim,* as referred to in Kant's categorical imperative, is an implied general principle or rule underlying a particular action. If, for example, I usually break my promises, then I act according to the private maxim that promise breaking is morally acceptable when it is in my best interests to do so. But can one take this maxim and transform it into a universal moral law? As a universal law this particular maxim would be expressed as follows: "It is permissible for everyone to break promises when it is in their best interests to do so." Such a law, however, is invalid because it entails both a pragmatic and a logical contradiction. There is a pragmatic (or practical) contradiction because the maxim is self-defeating if it is universalized. According to Korsgaard, "your action would become ineffectual for the achievement of your purpose if everyone (tried to) use it for that purpose."[22] Consider this example: An individual borrows some money from a friend and he promises to pay her back. However, he has no intention of keeping that promise. But this objective, that is, getting some money from her without repaying it, cannot be achieved by making a false promise in a world where this maxim has been universalized. As Korsgaard puts it, "The efficacy of the false promise as a means of securing money depends on the fact that not everyone uses promises that way."[23]

Universal promise breaking also implies a logical contradiction (such as a square circle); if everyone were to break their promises,

the entire institution of promising would collapse; there would be no such thing as a "promise" because in such a climate anyone making a promise would lack credibility. A world of universalized promise breaking is inconceivable. Thus, in view of the contradictions involved in universalizing promise breaking, we have a perfect duty to keep all of our promises.

Kant strongly implies that *perfect duties,* that is, duties that we are always obliged to follow, such as telling the truth or keeping a promise, entail both a logical and pragmatic contradiction. Violations of imperfect duties, however, are pragmatic contradictions. Korsgaard explains that "perfect duties of virtue arise because we must refrain from particular actions *against* humanity in our own person or that of another."[24] *Imperfect duties,* on the other hand, are duties to develop one's talents where the individual has the latitude to fulfill this duty using many different means.

Kant's categorical imperative is his ultimate ethical principle. It is the acid test of whether an action is right or wrong. According to Kant, then, any self-contradictory universalized maxims are morally forbidden. The categorical imperative functions as a guide, a "moral compass" that gives us a reliable way of determining a correct and consistent course of action. According to Norman Bowie, "the test of the categorical imperative becomes a principle of fair play—one of the essential features of fair play is that one should not make an exception of oneself."[25]

Also, from the categorical imperative we can derive other duties such as the duty to keep contracts, to tell the truth, to avoid injury to others, and so forth. Kant would maintain that each of these duties is also categorical, admitting of no exceptions, because the maxim underlying such an exception cannot be universalized.

How might we apply Kant's theory to the mundane ethical problems that arise in cyberspace? Consider the issue of intellectual property. As Korsgaard observes, "property is a practice,"[26] and this practice arguably makes sense for both physical property as well as intellectual property. But a maxim that permitted stealing of such property would be self-defeating. That maxim would say, "It's acceptable for me to steal the intellectual property validly owned by the creators or producers of that property." Such a universalized maxim, permitting everyone to take this intellectual property, is self-defeating precisely because it leads to the destruction of the entire "practice" of intellectual property protection. Because the maxim allowing an individual to freely appropriate another's intellectual property does not pass the universalization test, a moral

agent is acting immorally when he or she engages in acts such as the unauthorized copying of a digital movie or music file.[27]

At the heart of Kant's ethical system is the notion that there are rational constraints on what we can do. We may want to engage in some action (such as downloading copyrighted files), but we are inconsistent and hence unethical unless we accept the implications of everyone doing the same thing. According to Kant, it is unethical to make arbitrary exceptions for ourselves. In the simplest terms, the categorical imperative suggests the following question: What if everybody did what you are doing?

Before concluding this discussion on Kant, it is worth restating his second formulation of the categorical imperative: "Act in such a way that you treat humanity, whether in your own person or in the person of another, always at the same time as an end and never simply as a means."[28] For Kant as well as for other moralists (such as Finnis), the principle of humanity as an end in itself serves as a limiting condition of every person's freedom of action. We cannot exploit other human beings and treat them exclusively as a means to our ends or purposes. This could happen, for example, through actions that deceive one's fellow human beings or actions that force them to do things against their will. According to Korsgaard:

> According to [Kant's] Formula of Humanity, coercion and deception are the most fundamental forms of wrongdoing to others—the roots of all evil. Coercion and deception violate the conditions of possible assent, and all actions which depend for their nature and efficacy on their coercive or deceptive character are ones that others cannot assent to . . . Physical coercion treats someone's person as a tool; lying treats someone's reason as a tool.[29]

If we follow this categorical imperative, we will make sure that our projects and objectives do not supersede the worth of other human beings. This principle can also be summed up in the notion of *respect*. One way to express universal morality is in terms of the general principle of respect for other human beings who deserve that respect because of their dignity as free and rational persons.

One of the problems with Kant's moral philosophy is its rigidity. There are no exceptions to the moral laws derived from the absolute categorical imperative. Hence, lying is *always* wrong even though we can envision situations where telling a lie (e.g., to save a human life) is a reasonable and proper course of action. In cases such as this, there is a conflict of moral laws: the law to tell the truth and the law to save a

life in jeopardy, and we have no alternative but to admit an exception to one of them. As A. C. Ewing points out:

> In cases where two laws conflict it is hard to see how we can rationally decide between them except by considering the goodness or badness of the consequences. However important it is to tell the truth and however evil to lie, there are surely cases where much greater evils can still be averted by a lie, and is lying wrong then?[30]

Ewing's argument that it is difficult to avoid an appeal to consequences when two laws conflict poses problems for Kant's moral philosophy, despite its powerful appeal.

An alternative duty-based philosophy proposed by William D. Ross (1877–1940), a contemporary English philosopher, attempts to obviate the difficulties posed by Kant's inflexibility. Ross argues in his book *The Right and the Good*[31] that we are obliged to follow several basic *prima facie* duties that each of us can intuit through simple reflection. These duties are *prima facie* in the sense that they are conditional and not absolute. This means that under normal circumstances we must follow a particular duty, but in those unusual situations where duties conflict with one another, one duty may be overridden by another duty that is judged to be superior, at least under these specific circumstances. According to Ross, moral rules or principles are not categorical as they are for Kant, so they can have exceptions. Thus, a moral principle can be sacrificed or overridden, but only for another moral principle, not just for arbitrary, selfish, or even utilitarian reasons.

According to Ross, the seven *prima facie* moral duties that are binding on all moral agents are the following:

1. One ought to keep promises and tell the truth (*fidelity*).
2. One ought to right the wrongs that one has inflicted on others (*reparation*).
3. One ought to distribute goods justly (*justice*).
4. One ought to improve the lot of others with respect to virtue, intelligence, and happiness (*beneficence*).
5. One ought to improve oneself with respect to virtue and intelligence (*self-improvement*).
6. One ought to exhibit gratitude when appropriate (*gratitude*).
7. One ought to avoid injury to others (*noninjury*).

Ross makes little effort to provide any substantial rationalization or theoretical grounding of these duties. We might just say that they are common rules of morality, obvious to all rational humans because they have the general effect of reducing harm or evil to others.

The Achilles' heel of Ross's theory can be isolated by examining two specific problems: (1) his list of duties seems arbitrary because it is not metaphysically or even philosophically grounded, and (2) the list seems incomplete—where, for example, is the duty not to steal property from another? It may be included under the duty to avoid injury to others, but that is not altogether clear. Moreover, is it really true that all human beings (even those in different cultures) simply "intuit" these same principles? Finally, *The Right and the Good* provides little help for resolving situations where two *prima facie* duties do conflict. Ross offers few concrete criteria for determining when one obligation is more stringent and compelling than another.

Despite these shortcomings, however, Ross's framework, as with the others we have considered, is not without some merit. A focus on one's moral duty (or even conflicting duties) in a particular situation is a worthy starting point for moral reasoning about some dilemma or quandry. Further, for many moral conundrums, a sincere and rational person can develop sound, objective reasons for determining which duty should take priority.

New Natural Law

The natural law tradition has been neglected in most books on business and computer ethics. Detractors claim that it's too "impractical" and too closely associated with the theistic philosophy of St. Thomas Aquinas. MacIntyre, however, makes the case that the natural law ethic is superior to the "theories of those imprisoned within modernity [that] can provide only ideological rationalizations [such as] modern consequentialism and modern contractarianism."[32]

The new natural law, developed by John Finnis and Germain Grisez, remains faithful to the broad lines of natural law theory found in the philosophy of Aquinas. But it also attempts to make some necessary modifications demanded by the complexity of contemporary moral problems. Like Aquinas, Finnis and Grisez claim that the starting point of moral reflection is the first practical principle: "Good should be done and evil avoided," where *good* means what is intelligibly worthwhile. For the most part, human beings behave rationally and pursue what is good for them, what perfects their nature and makes them better off. But what is the good? Recall Finnis' argument that there are seven basic human goods that are the key to human flourishing: life and health, knowledge of the truth, play (and some forms of work), aesthetic experience, sociability (including friendship and marriage), religion, and practical reasonableness. All

of our choices ultimately point to one of these intelligible goods. For example, if someone asks Paul why he plays golf so much, he could answer that he enjoys the game or that he likes the exercise. The first answer points to the basic human good of play and the second to the good of health.

Each one of us participates in these basic goods, though we may participate in some goods more than others, and we do so to achieve "fullness of life." Practical reasonableness, which includes the value of authenticity, shapes one's participation in the other basic goods. And one requirement of practical reasonableness is that it is unreasonable to choose directly against any basic value, "whether in oneself or in one's fellow human beings."[33]

But how do we get from these basic human goods to specific moral norms and human rights? Our practical reason grasps that each of these basic human goods is an aspect of human flourishing and that a good in which any person shares also fulfills other persons. Whenever one intentionally destroys, impedes, or damages one of these goods that should be allowed to be, there is moral evil. Thus, we can stipulate the First Principle of Morality: *keep one's choices open to integral human fulfillment,* the fulfillment of all persons and communities.[34]

This principle, however, is too general and so we also need intermediate principles to specify the primary moral principle. Grisez calls these *modes of responsibility,* which include the Golden Rule (or the universalizability principle), "for a will marked by egoism or partiality cannot be open to integral human fulfillment."[35] These modes also include the imperative to avoid acting out of hostility or vengeance and never to choose evil as the means to a good end. The good or the end of my actions does not justify the use of unjust means that damage a basic good. According to this principle, for example, one could not justify telling a lie that damages the truth to advance a friendship. In this case, one is exercising favoritism with regard to these goods, which are incommensurable and all deserving of the same respect.

Specific moral norms can be deduced from those basic human goods with the help of the intermediate principles such as the Golden Rule. For example, because human life is a basic human good, certain acts such as the taking of innocent life are forbidden as a matter of natural law. Finnis states this natural law (or absolute moral norm) as follows: "Every act which is intended, whether as end or means, to kill an innocent human being and every act done by a private person which is intended to kill any human being" is prohibited.[36] This precludes necessary acts of self-defense. And from the basic good of knowledge of the truth, we can deduce the moral imperative of

veracity and "the right not to be positively lied to in any situation in which factual communication is reasonably expected."[37]

The new natural law provides a different vantage point from which to judge ethical conundrums in cyberspace. The value of this approach is its unwavering fidelity to the role of basic human goods such as life, health, and knowledge of the truth. It compels us to consider whether certain policies or actions are consistent with human flourishing, that is, with the realization of these basic human goods identified by Finnis and Grisez. It is difficult to argue, for instance, that deceptive spamming has any moral legitimacy; by undermining the truth in factual Internet communications, this form of spam deserves to be classified as morally reprehensible. The natural-law framework allows us to appreciate why this is so wrong by focusing on its true negative impact.

Although Finnis and Grisez have tried to disengage the natural-law framework from the metaphysics of Aquinas, critics claim that they do not succeed. According to Lisska, "One intuits the basic goods and it just happens that set of goods correspond to human well being. But what establishes the causal relationship?"[38] Nonetheless, according to Grisez, this theory attempts to combine the strengths of teleology and deontology. It grounds morality in human goods, "the goods of real people living in the world of experience," and it protects each person's dignity with intermediate principles and moral absolutes.[39]

Postscript on Moral Theory

As we have seen, none of these theories are without flaws or contradictions, but they do represent viable avenues for reasoning about moral issues, especially when those issues go beyond the level of moral common sense. They also have certain elements in common, particularly an orientation to "the other"—along with the need to consider the interests and perspectives of the affected parties in assessing alternative action plans, the other's moral and legal rights, and our duty to treat the other as an end and not as a means. And they all stand in opposition to the dangerous and myopic philosophy of ethical egoism, which is blind to the rights and aspirations of others.

Before concluding this material on ethical theories, we can summarize how they can be applied to some of the moral quandaries that arise in the electronic frontier of cyberspace. **TABLE 1-2** provides a concise framework for putting these four basic theories into action.

In some cases these four frameworks converge on the same solution to an ethical quandry. At other times, they suggest different solutions

TABLE 1-2 Summary of Ethical Frameworks

Theory Type	Operative Questions
Consequentialism/utilitarianism	Which action or policy generates the best overall consequences or the greatest net expectable utility for all affected parties?
Duty-based morality	Can the maxim underlying the course of action being considered be universalized? Is the principle of fair play being violated? If there appears to be conflicting duties, which is the stronger duty?
Rights-based morality	Which action or policy best protects the human and legal rights of the individuals involved?
New natural law	Does the proposed action or policy promote the basic requirements of human flourishing? Does it impede, damage, or destroy any of the basic human goods?

to the problem and one must decide which framework should "trump" or override the others. Should one respect the rights of some group or individual, even though following that alternative will be less beneficial to all affected parties than other alternatives? Resolving such questions requires careful and objective reasoning, but responsible behavior sometimes requires that this extra step be taken. To be sure, the Internet presents unique ethical challenges that could never have been envisioned by Aquinas, Kant, or Mill, but these frameworks still provide a general way of coming to terms with these tough questions.

Floridi's *Macroethics*

Before concluding this discussion, it is worth considering a new high-level theory specifically designed to accommodate our contemporary Information Age, which is so irreversibly centered on digital information. Despite the breadth and depth of traditional ethical theories, some contemporary philosophers believe that they are inadequate to address the complex moral problems of our networked information society. One such thinker is Luciano Floridi, who finds fault with these traditional approaches because they are too anthropocentric or too preoccupied with how personal actions affect other persons. Those theories

pay little attention to how actions impact the broader biological, social, and informational environment. As a complement to those theories, Floridi proposes his more ecological macroethics, or Information Ethics (IE).

Floridi's ethical theory has three major characteristics: it is ontocentric, ecological, and patient-oriented. First, what does he mean by "ontocentric"? At the core of Floridi's theory is the thesis that all entities in the universe, both animate and inanimate, are informational objects or "clusters of data," and this common feature endows them with some moral value. This category of beings deserving moral consideration includes even digital objects that exist only in cyberspace or in a database because they, too, are obviously informational objects. As a result, ethical discourse and reasoning must take into account the moral status of all entities in the infosphere. Floridi explains that according to IE, "even ideal, intangible, or intellectual objects can have a minimal degree of moral value."[40]

Although biocentrists maintain that we should not needlessly destroy or harm any living being, the ontocentrist espouses the belief that no being or informational object should be damaged or destroyed by the alteration of that being's data structure without sufficient reason. Being, therefore, is more fundamental than life. According to Floridi, all beings have the Spinozian right to persist in being and a "constructionist right to flourish."[41] Of course, the moral worth of certain informational objects is minimal and "overrideable," but even these objects still warrant some degree of moral consideration. Ontocentrism, Florid maintains, is the only authentic ecology because of its sensitivity to the entire infosphere.

IE is a "patient-oriented" theory because it is concerned with what qualifies as a moral patient, that is, an object worthy of moral consideration. Because all information objects qua information objects have intrinsic value, they qualify as moral patients, worthy of some degree of moral worth. In this moral framework, evil is equated with entropy, which refers to any kind of "disruption, corruption, pollution, and depletion of informational objects."[42] Floridi's chief concern is the welfare of the whole infosphere. IE is a macroethics precisely because of its interest in the entire infosphere and the entropy or impoverishment of being that could happen to any entity that occupies this environment.

Floridi's theory is also concerned with the theme of moral agency, and once again he departs from the anthropocentric assumptions of traditional ethical theory. Floridi broadens the class of moral agents to include robots, software bots, and other information technology (IT) systems. He defines the moral agent as an interactive, autonomous, and adaptable transition system capable of performing "morally

qualifiable" actions, that is, actions that can cause good or evil. A transition system is one that changes its states, and this system is interactive when it acts upon and is affected by the environment. That system is autonomous when it can change its state without direct response to interaction, and it is adaptable when those interactions change the transition rules. Given these criteria, we can reasonably conclude that artificial agents like robots have some degree of moral agency. Floridi concedes that although artificial moral agents occupying the infosphere, such as robots and corporations, can be held morally accountable, they lack moral responsibility for their actions. In the infosphere, however, we must transition from a responsibility-oriented ethics based on punishment and reward to an ethics based on "accountability and censure."[43]

In this text we only tangentially explore the role of artifacts in cyberspace such as surveillance tools and software bots that collect information for search engines and other data aggregators. The reader might ponder whether these entities have any sort of artificial moral agency, if considered from Floridi's nonanthropocentric perspective. Also, as these artifacts become more sophisticated and "intelligent," the debate about their moral status will surely intensify.

As with the other theories we have considered, thoughtful critics point to certain shortcomings. They question the premises of ontocentrism, which assumes that every being, including a rock or a piece of spam email, has some degree of moral worth. Others argue that this abstract theory is not as useful or broadly applicable as utilitarianism or rights-based approaches to ethics. Floridi insists that IE is not meant as a substitute for traditional ethics but as a supplement. He admits, however, that we need "an ethical framework that can treat the infosphere as a new environment worth the moral attention and care of the human inforgs inhabiting it."[44]

Normative Principles

Those who find ethical theory too abstract can turn to an approach known as *principlism*. It is commonly used in biomedical ethics and has become popularized through the work of Beauchamp and Childress.[45] These moral principles are derived from and are compatible with all of the moral theories articulated here. They constitute *prima facie* duties that are always in force but may conflict on occasion. The four principles proposed by Beauchamp and Childress are autonomy, nonmaleficence, beneficence, and justice. Those who advocate

this approach also prescribe certain "prudential requirements" that determine when one *prima facie* principle should be given more weight than another. These include "being sure that there is a realistic prospect of achieving the moral objective one has chosen to honor; no alternative course of action is possible that would honor both conflicting obligations; and we minimize the effects of infringing on the prima facie duty."[46] A brief sketch of these four principles follows.

Autonomy

Kant and other philosophers have consistently argued that a defining element of personhood is one's capacity to be autonomous or self-determining. According to Gary Doppelt, "the Kantian conception of personhood ties the moral identity of persons to the supreme value of their rational capacities for normative self-determination."[47] All rational persons have two key moral powers or capacities: they possess the ability to develop and revise a rational plan to pursue their conception of the good life, and they possess the capacity to respect this same capacity of self-determination in others. Thus, autonomy is not only a necessary condition of moral responsibility, it is also through the exercise of autonomy that individuals shape their destiny according to their notion of the best sort of life worth living. When someone is deprived of their autonomy, their plans are interfered with and they are not treated with the respect they deserve. Of course, respect for autonomy must be balanced against other moral considerations and claims.

Nonmaleficence

The principle of nonmaleficence can best be summarized in the moral injunction: "Above all, do no harm." According to this core principle, one ought to avoid unnecessary harm or injury to others whenever possible. This negative injunction against doing injury to others is sometimes called the "moral minimum." However one may choose to develop a moral code of conduct, this injunction must be given a preeminent status. Most moral systems go well beyond this minimum requirement, as we have seen in the theories already discussed, but that does not detract from the central importance of this principle. According to Jon Gunneman and his coauthors,

> We know of no societies, from the literature of anthropology or comparative ethics, whose moral codes do not contain some injunction against harming others. The specific notion of *harm* or *social injury* may vary, as well as the mode of correction and restitution but the injunctions are present.[48]

Beneficence

This is a positive duty and has been formulated in many ways. In the simplest terms it means that we should act in such a way that we advance the welfare of other people when we are able to do so. In other words, we have a duty to help others. But what does this really mean? When am I duty bound to help another person or even an institution? It is obvious that we cannot help everyone or intervene in every situation when someone is in need. Hence, some criteria are necessary for determining when such a moral obligation arises. In general, it can be argued that we have a duty to help others under the following conditions:

1. The need is serious or urgent.
2. We have knowledge or awareness of the situation.
3. We have the capability to provide assistance ("ought assumes can" is the operative principle).

If, for instance, one is an Olympic swimmer and sees someone drowning at the beach, one has an obligation to attempt a rescue of that person, especially if this is the only recourse and there is little risk to one's own life. This principle has some relevance when we evaluate society's questionable duty of beneficence to provide universal Internet service.

Justice

Although theories of justice have their differences, most have a common adherence to this basic formal principle: "Similar cases ought to be treated in similar ways." Above all else, justice requires fair treatment and impartiality. This is a formal procedural principle of justice and needs to be supplemented by the criteria for determining "similar" cases. This leads into theories of distributive justice, which attempt to formulate an underlying principle for how we should distribute the benefits and burdens of social life. Some theories emphasize equality, that is, all goods should be distributed equally. John Rawls, for example, adopts an egalitarian approach, though he does argue that an unequal distribution of goods is acceptable when it works for the advantage of everyone, especially the least advantaged (the difference principle).[49] Other theories emphasize contribution and effort as formulated in this maxim: "Benefits or resources should be distributed according to the contribution each individual makes to the furtherance of society's goals." And still another theory of justice that has typically been associated with socialism argues for justice based on need: "From each according to his ability, to each according to his needs."[50]

Our purpose here is not to defend one of these theories against the other, but to illustrate that moral judgements should be based in part on the formal principle of justice and take into account some standard regarding how the benefits and burdens should be fairly distributed within a group or society at large.

There is no reason that these formal moral principles cannot be applied to some of the controversial problems that we consider in this text. They are certainly general enough to have applicability in the field of computer and Internet ethics as well as bioethics. A person who makes choices and develops policies attentive to the core human goods and to these more practical principles that generally promote those goods would surely be acting with the care and prudence that is consistent with the moral point of view.

DISCUSSION QUESTIONS

1. Do you agree with the philosophy of technological realism?
2. Explain the basic elements of Lessig's framework. What does he mean when he says that in cyberspace "the code is the law"?
3. Explain and critically analyze the essentials of Kant's moral theory.
4. In your estimation, which of the moral frameworks presented in this chapter has the most promise for dealing with the moral dilemmas that arise in cyberspace?

REFERENCES

1. Aristotle, *Nicomachean Ethics,* I, 3: 1095a6.
2. Larry Lessig, *Code and Other Laws of Cyberspace* (New York: Basic Books, 1999), 236.
3. See Frank Easterbrook, "Cyberspace and the Law of the Horse," *University of Chicago Law Forum* 207 (1996).
4. Phillippa Foot, "Moral Relativism" (1979 Lindley Lecture, Department of Philosophy, University of Kansas).
5. Larry Lessig, "The Laws of Cyberspace;" available at: http://cyberlaw.stanford.edu/lessig
6. Larry Lessig, "Tyranny in the Infrastructure," *Wired* 5.07 (1997): 96.
7. Jim Moor, "Just Consequentialism and Computing," in *Readings in Cyberethics,* ed. Richard Spinello and Herman Tavani (Sudbury, MA: Jones and Bartlett, 2001), 100.
8. Immanuel Kant, *Foundations of the Metaphysics of Morals* (Indianapolis: Bobbs Merrill, 1959), 33.
9. John Finnis, *Fundamentals of Ethics* (Washington, D.C.: Georgetown University Press, 1983), 125.
10. Jacques Ellul, *The Technological Society,* trans. John Wilkinson (New York: Vintage Books, 1964), xxv.
11. Ibid., 14.
12. Langdon Winner, *Autonomous Technology: Technics-out-of Control as a Theme of Political Thought* (Cambridge: MIT Press, 1977), 229.
13. Priscilla Regan, *Legislating Privacy* (Chapel Hill: University of North Carolina Press, 1995), 12.

14. Charles Taylor, *The Ethics of Authenticity* (Cambridge: Harvard University Press, 1991), 101.
15. Kenneth Goodpaster, "Some Avenues for Ethical Analysis in Management," in *Policies and Persons*, eds. John Matthews et al. (New York: McGraw-Hill, 1985), 495.
16. See, for example, James Rachels, *The Elements of Moral Philosophy* (New York: Random House, 1986); and William K. Frankena, *Ethics* (Englewood Cliffs, NJ: Prentice-Hall, 1963).
17. Frankena, *Ethics,* 29.
18. Charles Kelbley, "Freedom from the Good," in *Freedom and Value,* ed. Robert Johann (New York: Fordham University Press, 1975), 173.
19. Goodpaster, 497.
20. Kant, 16.
21. Kant, 18.
22. Christine Korsgaard, *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996), 78.
23. Ibid, 92.
24. Ibid, 21.
25. Norman Bowie, *Business Ethics: A Kantian Perspective* (Oxford: Blackwell Publishers, 1999), 17.
26. See Korsgaard, 97.
27. Richard A. Spinello, "Beyond Copyright: A Moral Investigation of Intellectual Property Protection in Cyberspace," in *The Impact of the Internet on Our Moral Lives,* ed. Robert J. Cavalier (Albany, NY: SUNY Press, 2005), 27–48.
28. Kant, 36.
29. Korsgaard, 194.
30. A. C. Ewing, *Ethics* (New York: Free Press, 1965), 58.
31. William D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930).
32. Alasdair MacIntyre, *Three Rival Versions of Moral Enquiry: Encyclopaedia, Genealogy, Tradition* (Notre Dame, IN: University of Notre Dame Press, 1990), 194.
33. John Finnis, *Natural Law and Natural Rights* (New York: Oxford University Press, 1980), 225.
34. Germain Grisez, "A Contemporary Natural Law Ethic," in *Normative Ethics and Objective Reason,* ed. George F. McLean (2001). http://www.mccc.edu/~howarthk/Grisez_Ethical_Theory.docx.
35. Ibid.
36. John Finnis, *Aquinas* (Oxford: Oxford University Press, 1998), 141.
37. Finnis, *Natural Law and Natural Rights,* 197.
38. Anthony Lisska, *Aquinas' Theory of Natural Law* (New York: Oxford University Press, 1996), 161.
39. Grisez.
40. Luciano Floridi, "Information Ethics," in *The Cambridge Handbook of Information and Computer Ethics,* ed. Luciano Floridi (Cambridge, UK: Cambridge University Press, 2010), 85.
41. Ibid., 84
42. Ibid.
43. Ibid., 88.
44. Luciano Floridi, "Ethics after the Revolution," in *The Cambridge Handbook of Information and Computer Ethics,* ed. Luciano Floridi (Cambridge, UK: Cambridge University Press, 2010), 19.

45. Thomas Beauchamp and J. F. Childress, *Principles of Biomedical Ethics,* 4th ed. (New York: Oxford University Press, 1994).
46. Mark Kaczeski, "Casuistry and the Four Principles Approach," in *Encyclopedia of Applied Ethics,* vol. 1, ed. Ruth Chadwick (San Diego, CA: Academic Press, 1998), 430.
47. Gary Doppelt, "Beyond liberalism and communitarianism: A critical theory of social justice," *Philosophy and Social Criticism* 14 (1988): 278.
48. Jon Gunneman et al., *The Ethical Investor* (New Haven, CT: Yale University Press, 1972), 20.
49. John Rawls, *A Theory of Justice* (Cambridge: Harvard University Press, 1971), 85–90.
50. Karl Marx, *Critique of the Gotha Program* (London: Lawrence and Werhart, Ltd., 1938), 14.

ADDITIONAL RESOURCES

Baase, Sara. *A Gift of Fire*: *Social, Legal and Ethical Issues in Computing*. 3rd ed. Upper Saddle River, NJ: Prentice-Hall, 2013.

Bynum, Terrell Ward. *Information Ethics: An Introduction*. Cambridge, MA: Blackwell, 1998.

Bynum, Terrell Ward, and Simon Rogerson. *Computer Ethics and Professional Responsibility*. London: Blackwell, 2002.

Cavalier, Robert, ed. *The Impact of the Internet on Our Moral Lives*. Albany, SUNY Press, 2005.

Collste, Goran, ed. *Ethics in the Age of Information Technology*. Linkoping, Sweden: Linkopings Universitet Centre for Applied Ethics, 2000.

Edgar, Stacey. *Morality and Machines*. 2nd ed. Sudbury, MA: Jones and Bartlett, 2003.

Ermann, M. David, Mary Williams, and Michele Shauf, eds. *Computers, Ethics, and Society*. 2nd ed. New York: Oxford University Press, 1997.

Floridi, Luciano, ed. *The Cambridge Handbook of Information and Computer Ethics*. Cambridge: Cambridge University Press, 2010.

Floridi, Luciano, *Philosophy and Computing: An Introduction*. London: Routledge, 1999.

Hester, D. Micah, and Paul Ford, eds. *Computers and Ethics in the Cyberage*. Upper Saddle River, NJ: Prentice-Hall, 2001.

Himma, Kenneth, and Herman Tavani, eds. *The Handbook of Information and Computer Ethics*. Hoboken, NJ: Wiley, 2008.

Johnson, Deborah. *Computer Ethics*. 4th ed. Upper Saddle River, NJ: Prentice-Hall, 2009.

Johnson, Deborah, and Helen Nissenbaum, eds. *Computers, Ethics and Social Values*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

Kling, Rob, ed. *Computerization and Controversy*. 2nd ed. San Diego: Academic Press, Inc., 1996.

Moore, Adam, ed. *Information Ethics*. Seattle: University of Washington Press, 2005.

Rosenberg, Richard. *The Social Impact of Computers*. 2nd ed. San Diego: Academic Press, 1997.

Spinello, Richard. *Case Studies in Information Technology Ethics*. Upper Saddle River, NJ: Prentice-Hall, 2002.

Spinello, Richard, and Herman Tavani, eds. *Readings in Cyberethics*. 2nd ed. Sudbury, MA: Jones and Bartlett, 2004.

Tavani, Herman. "The State of Computer Ethics as a Philosophical Field of Inquiry." *Ethics and Information Technology* 3, no. 2 (2001): 97–108.

Tavani, Herman. *Ethics and Technology: Controversies, Questions and Strategies for Ethical Computing*. New York: Wiley, 2012.