

7

EVALUATING AND CONTROLLING TECHNOLOGY

7.1 Evaluating Information

7.2 The “Digital Divide”

7.3 Neo-Luddite Views of Computers, Technology, and Quality of Life

7.4 Making Decisions About Technology

Exercises



In this chapter we consider such questions as these: Does the openness and “democracy” of the Web increase distribution of useful information or inaccurate, foolish, and biased information? How can we evaluate complex computer models of physical and social phenomena? How does access to digital technology differ among different populations? Is computing technology evil? Why do some people think it is? How should we control technology to ensure positive uses and consequences? How soon will robots be more intelligent than people? What will happen after that?

Whole books focus on these topics. The presentations here are necessarily brief. They introduce some of the issues and arguments.

7.1 Evaluating Information

*A little learning is a dang'rous thing;
Drink deep, or taste not the Pierian spring;
There shallow draughts intoxicate the brain,
And drinking largely sobers us again.*

—Alexander Pope, 1709¹

7.1.1 THE NEED FOR RESPONSIBLE JUDGMENT

Expert information or the “wisdom of the crowd”?

*We can get the wrong answer to a question quicker than our fathers
and mothers could find a pencil.*

—Robert McHenry²

There is a daunting amount of information on the Web—and much of it is wrong. Quack medical cures abound. Distorted history, errors, outdated information, bad financial advice—it is all there. Marketers and public relations firms spread unlabeled advertisements through blogs, social media, and video sites. Search engines have largely replaced librarians for finding information, but search engines rank Web pages by popularity (at least partly) and give prominent display to content providers who pay them; librarians do not. Wikipedia, the biggest online encyclopedia, is immensely popular, but can we rely on its accuracy and objectivity when anyone can edit any article at any time? On social journalism sites, readers submit and vote on news stories. Is this a good way to get news? The nature of the Internet encourages people to post their immediate thoughts and reactions without taking time for contemplation or for checking facts. How do we know what is worth reading in contexts where there are no editors selecting the well-written and well-researched?

When we see a video of a currently popular performer singing with Elvis Presley, we know we are watching creative entertainment—digital magic at work. But the same technology can deceive. Video-manipulation tools (and increased bandwidth) provide the opportunity for “forging” people. A company developed an animation system that modifies video images of a real person to produce a new video in which the person is speaking whatever words the user of the system provides. Another system analyzes recordings of a person’s voice and synthesizes speech with the voice, inflections, and tones of that person. Combined, these systems will likely have many uses, including entertainment and advertising, but clearly people can also use them to mislead in highly unethical ways.³ How do we know when someone is manipulating us?

Example: Wikipedia

To explore some issues of information quality, we consider Wikipedia. Wikipedia is a collaborative project among large numbers of strangers worldwide. It is huge, free, participatory, noncommercial, ad-free, and written by volunteers. The English edition has almost four million articles, more than 10 times as many as the long-respected Encyclopaedia Britannica, first published in 1768 and online since 1994.⁴ Wikipedia is one of the Internet’s most-used reference sites. But are its entries true, honest, and reliable?

We expect encyclopedias to be accurate and objective. Traditionally, expert scholars selected by editorial boards write encyclopedias. Volunteers, not carefully selected scholars, write and continually edit and update Wikipedia articles. Anyone who chooses to participate can do so. People worry that the lack of editorial control means no accountability, no standards of quality, no way for the ordinary person to judge the value of the information. They argue that because hundreds of millions of people—anyone at all—can write or edit articles, accuracy and quality are impossible. Truth does not come from populist free-for-alls. Members of the staffs of political candidates have distorted the Wikipedia biographies of their candidates to make their bosses look better. Opponents and enemies regularly vandalize profiles of prominent people. The staff of a federal agency removed criticisms of the agency from its Wikipedia article. Discredited theories about historic events such as the terrorist attacks on September 11, 2001, and the assassination of John F. Kennedy reappear regularly. A lawyer reported that one party in a case edited Wikipedia entries to make information appear more favorable to that party. (Jurors are not supposed to consult online sources about a trial, but some do.) Removing false information, hoaxes, and the like requires constant effort, according to Wikipedia volunteers. The Encyclopaedia Britannica has had errors and oddities, but the nature of Wikipedia makes it prone to more. Anonymity of writers encourages dishonesty. Open, volunteer, instant-publishing systems cannot prevent errors and vandalism as easily as publishers of printed books or closed, proprietary online information sources.

In spite of the errors, sloppiness, bad writing, and intentional distortions, most of Wikipedia is, perhaps surprisingly, of high quality and extraordinary value. Why? What

protects quality in large, open, volunteer projects? First, although anyone *can* write and edit Wikipedia articles, most people do not. Thousands write and edit regularly, not millions. Most are educated and have expertise in the subjects they write about. They correct articles promptly. (Wikipedia saves old versions, so it can restore an article someone has vandalized.) After well-publicized incidents of manipulation of articles, Wikipedia's managers developed procedures and policies to reduce the likelihood of such incidents. For example, they lock articles on some controversial topics or people; the public cannot directly edit them.

We, as users, can (and must) learn to deal appropriately with side effects or weaknesses of new paradigms. Even though so much of Wikipedia is excellent and useful, we learn that someone might have wrecked the accuracy and objectivity of any individual article at any hour. We learn that articles on technology, basic science, history, and literature are more likely to be reliable than those on politics, controversial topics and people, and current events. We learn to use Wikipedia for background, but to check facts and alternative points of view. Should we judge Wikipedia (and, by extension, the mass of information on the Web) by the excellent material it provides or by the poor-quality material it includes?

Written by fools for the reading of imbeciles.

—An evaluation of newspapers, not websites, by a character in Joseph Conrad's novel *The Secret Agent* (1907)

The “wisdom of the crowd”

People ask all sorts of questions on Yahoo! Answers (and other sites like it) about dating, make-up, food, college (“Are online college classes as good as classroom classes?”), and wide-ranging technical, social, economic, and political issues (“If we can produce enough food to feed everyone in the world, why don't we?”) Of course, a lot of answers are ill-informed. The questioner designates the posted answer he or she deems the best. What qualifies the questioner, presumably a person who does not know the answer, to judge the worthiness of the replies? To what extent does the ease of posting a question reduce the likelihood that a person will seek out well-researched or expert information on the subject? There are obviously questions for which this kind of forum might not provide the best results. (An example might be: Is it safe to drink alcohol while using an acne medicine?) However, the first two sample questions I quoted above are likely to generate a lot of ideas and perspectives. Sometimes that is exactly what the questioner wants. Without the Web, if someone asked questions like those of only a few friends, the answers might be less varied and less useful.

Some health sites on the Web encourage the public to rate doctors, hospitals, and medical treatments. Are such ratings valuable or dangerous? Will they motivate doctors and hospitals to change their practices to achieve higher ratings at the expense of good medical care? Steve Case, co-founder of AOL and founder of a health site that emphasizes ratings by the public, argues that if millions of people participate, the results will be very useful. Others are extremely suspicious of “the wisdom of the crowd.” And there is always

concern about manipulation. Websites have sprung up to buy and sell votes to get prominent display for articles on social media sites. What are the implications of such practices for sites where the public rates medical care? Will providers of new or questionable medical treatments generate fake favorable reviews and votes? Will responsible operators of sites that display material based on rankings or votes anticipate manipulation and protect against it?

Let's pause briefly to put the problems of incorrect, distorted, and manipulated information in perspective. Quack medical cures and manipulative marketing are hardly new. Product promotions not labeled as advertising date back hundreds of years. Eighteenth-century opera stars paid people to attend performances and cheer for them or boo their rivals. "Hatchet jobs" in the form of news articles, books, ads, and campaign flyers have dishonestly attacked politicians long before the Web existed. There are plenty of poorly written and inaccurate books. Historical movies merge truth and fiction, some for dramatic purposes, some for ideological purposes. They leave us with a distorted idea of what really happened. Two hundred years ago, cities had many more newspapers than they do today. Most were opinionated and partisan. At supermarket counters, we can buy newspapers with stories as outlandish as any online. The *New York Times* is a prime example of a respected newspaper, staffed by trained journalists, with an editorial board in charge. Yet one of its reporters fabricated many stories. Numerous other incidents of plagiarism, fabrication, and insufficient fact-checking have embarrassed newspapers and television networks.

OK, the problems of unreliable information are not new. But they are problems, and the Web magnifies them. So we consider two questions: How good is the wisdom of the crowd? And how can we distinguish good sources of information on the Web?

Researchers find that crowds do, in fact, generate good answers to certain kinds of questions. When a large number of people respond, they produce a lot of answers, but the average, or median, or most common answer is often a good one. This works well when the people are isolated from each other and express independent opinions. Some researchers think a large (independent) group is likely to be more accurate than a committee of experts for a variety of questions such as estimating economic growth or how well a new product or movie will do. (A Canadian mining company, perhaps hoping for such a phenomenon, posted a large set of geological data on the Web and held a contest to choose areas to look for gold.) However, when people see the responses provided by others, some undesirable things happen. People modify their responses so that the set of responses becomes less diverse, and the best answer may no longer stand out. People become more confident from reinforcement even though accuracy has not improved. The wisdom of crowds depends on diversity and independence. In social networks (as well as in-person teams working on projects in businesses, organizations, and government agencies), peer pressure and dominant personalities can reduce the wisdom of the group.⁵

How can we distinguish good sources of information on the Web? Search engines and other services at first ranked sites by the number of people who visit them. Some developed more sophisticated algorithms to consider the quality of information on sites

where users provide content. (In response, some sites added editors and fact-checking to improve quality.) A variety of people and services review and rate sites and blogs. Critics of the quality of information on the Web and the lack of editorial control disdain such ratings as merely popularity contests, contending, for example, that the Internet gratifies the “mediocrity of the masses.”⁶ For blogs, as for Wikipedia or health care sites, they argue that popularity, voting, and consensus do not determine truth. That is correct, but there is no magic formula that tells us what is true and reliable either on the Web or off the Web. That a large number of people visit a website does not guarantee quality, but it provides some information. (Why have newspapers long published “best seller” lists for books?) We can choose to read only blogs written by Nobel Prize winners and college professors, if we wish, or only those recommended by friends and others we trust. We can choose to read only product reviews written by professionals, or we can read reviews posted by the public and get an overview of different points of view.

Over time, the distinction between the online equivalents of responsible journalism and supermarket tabloids becomes clear. Good reputations develop, just as they have for decades offline. Many university libraries provide guides for evaluating websites and the information on them. (I list some at the end of this chapter.) One good step is to determine who sponsors the site. If you cannot determine the sponsor of a site, you can consider its information as reliable as the information on a flyer you might find under your car’s windshield wiper when you park in a busy parking lot. Ultimately, we must find sites, reviewers, ratings, editors, experts, and other sources we trust. Good judgment and skepticism are always useful.

The only way to preserve the wisdom of the crowd is to protect the independence of the individual.

—Jonah Lehrer⁷

Vulnerable viewers

Since you are reading this book, you probably are a student, a reasonably well-educated person who is learning how to analyze arguments and make good judgments. You can develop skills to evaluate material you read on the Web. But what about people who have less education or ability? For example, what risks does bad information pose to children who find it on the Web? Some critics of the Web worry most about the impact of inaccurate information on such vulnerable people. The fears of some seem to edge toward a belief that we (or experts, or the government) should somehow prevent such information from appearing. The many strong arguments for freedom of speech in general are arguments against any centralized or legally mandated way of accomplishing this. What can we do to improve the quality of information? Basic social and legal forces help (to a degree): freedom of speech (to provide responses, corrections, alternative viewpoints, and so on), teachers and parents, competition, fraud and libel laws—and people who care,

who volunteer to write, review, and correct online information. What else can we do to reduce access to dangerously wrong information by vulnerable people?

Narrowing the information stream

All the problems of junk and nonsense on the Web notwithstanding, the Web now gives us access to more high-quality, up-to-date information than libraries did in the past. Consider current events, politics, and controversial issues. We can read and listen to thousands of news sources on the Web from our own and other countries, getting different cultural and political perspectives on events. We can read the full text of government documents—bills, budgets, investigative reports, congressional testimony and debate—instead of relying on a few sentences quoted from an official news release or a sound bite from a biased spokesperson. We can search archives of millions of news articles from the past 200 years. We can follow websites, blogs, tweets, and social media news of conservatives, liberals, libertarians, tea party activists, environmentalists, evangelical Christians, animal rights activists, and so on, far more easily and cheaply than when we had to seek out and subscribe to their print newsletters and magazines. But what do people actually do? Some get all their news and interpretation of events from a small number of sites that reflect a specific political point of view. Online tools make it easy: you just set up your bookmarks and feeds and never look anywhere else, except at other sites recommended by the ones you frequent. Some critics see the Web as significantly encouraging political narrowness and political extremes by making it easy for people to avoid seeing alternative opinions.

The phenomenon of using the information that is easy to get applies to other fields besides politics, of course. I hear sad complaints from librarians and experienced researchers: Too many students and professionals believe whatever they read in Wikipedia. Researchers “reinvent the wheel” (or apply for grants to do so) because they do not read relevant work in their field published in the past and available only in paper journals. Because there is so much on the Web, it is too easy to ignore what is not there. How serious is the problem of ignoring older, undigitized work? Is it a temporary problem that will go away when virtually all research is available electronically?

A researcher analyzed millions of academic articles published over 50 years and found that as journals moved online, authors tend to cite fewer articles, more recent ones, and articles from a narrower set. The speculation is that researchers using search engines to find articles related to their work select from among the ones that appear high in search results—the ones that are already cited frequently. Those articles might indeed be the most important, but this approach reinforces previous choices and can lead researchers to miss less popular but very relevant work. Researchers have far more (and easier) access to articles and journals online than they had in the stacks of libraries. However, as the author of the study says, searching online “puts researchers in touch with prevailing opinions, but this may accelerate consensus and narrow the range of findings and ideas built upon.”⁸ The effect of accelerating consensus and narrowing results is similar to what researchers saw with the wisdom of the crowd when crowd members were not

Idiots and dunderheads

A fool and his money are soon parted.
—Old English proverb

New technologies can have the unintended side effect of diminishing older skills. Computing technology has reduced the use of cursive writing, for example. Microsoft made a conscious decision with the effect of diminishing language skills. The thesaurus in Microsoft Word 2000 (and some later versions) lists the verb “trick” as the only meaning for “fool.” It omits noun synonyms “clown,” “blockhead,” “idiot,” “ninny,” “dunderhead,” “ignoramus,” and others—all present in earlier versions. Standard

references such as dictionaries and Roget’s Thesaurus contain some of these and more choices.

Microsoft said it eliminated words “that may have offensive uses.”^{9*} Was this a dunderheaded decision that dulls the language and reduces literacy? Do producers of widely used reference works have an ethical responsibility to report the substance of their field accurately, or a social responsibility to remove potentially offensive words from the language?

* Microsoft restored some synonyms meaning a foolish person but continues to omit the more colorful and more offensive terms.

independent, though the mechanism is different. Clearly, it is good for researchers to be aware of this phenomenon and to broaden their searches when appropriate. The number of scholarly papers published each year has grown enormously (to, very roughly, a million yearly). Is it the tendency to use search tools in a somewhat lazy way—or the sheer number of papers—that causes some valuable work to be missed?

If we receive too much information that does not interest us, we stop reading it. To counter this problem, Facebook implemented algorithms to filter news feed updates from friends based on how recently a member communicated with them. But of course, sometimes we want to hear from those people we have not heard from in a long time. What better methods could Facebook use? And is this relevant to social issues beyond personal relationships? It is. Eli Pariser, president of (liberal) MoveOn.org, includes conservatives among his Facebook friends because he wants to be aware of views different from his own. Over time, he realized he was no longer receiving updates from them (because he did not communicate with them regularly). Although Facebook members can turn off the filtering of news feeds, most people are not aware of it. Pariser considered the problem of filtered information so disturbing that he wrote a book about it.¹⁰ What lessons can we learn from Facebook’s filtering? It is not ideal to use Facebook as our main source of access to political discussion. Facebook’s choice of a default setting (filtering turned on) might not be best (but, then again, most people might prefer it). More fundamentally, the problem of determining what information is relevant and desired does not have an easy or obvious solution. We observed (in the box in Section 3.2.4) that in some situations filtering out too much is better than filtering out too little, while in other situations the opposite is true. Any solution that a search engine or social media service adopts will not

be perfect. However, when people do not know that they are seeing filtered information, they do not know to turn off the filter or to look elsewhere for more information. There should be a clear indication when filters are active and an explanation of what they do.

How else does the Web narrow information streams? In Chapter 2 we saw that search engines personalize results for users based on their location, past searches, profile information, and other criteria. Given the huge amount of information on the Web, this fine tuning helps us find what we want quickly. It is very valuable. However, it does mean that when we are searching for something outside our usual context, including perhaps information on controversial subjects, we might have to make an effort to look a little harder.

Do the various aspects of the Internet that narrow our information stream significantly diminish access to different points of view on controversial social and political topics? Members of radical political groups (left and right) and cults got information and opinions from narrow sources well before the Web. Does the Web encourage or increase ideological isolation? Does it simply reflect the choices that some people have often made in the past? Does it make it more likely that we will see a variety of points of view? When we criticize aspects of the Web, it is helpful to look to human nature and the past for perspective. It is also helpful to look toward an ideal to suggest improvements.


Abdicating responsibility

*I have a spelling checker.
It came with my PC.
It plainly marks four my revue,
Miss steaks aye can knot sea.
Eye ran this poem threw it,
I'm sure your pleased too no.
It's letter perfect in it's weigh,
My checker tolled me sew.*

—Jerrold H. Zar, “Candidate for a Pullet Surprise”¹¹

The tools and technologies we use encourage certain practices and consequences by making them easier. The spelling-checker verse above humorously illustrates the problem of doing what the tool makes easy and ignoring other important tasks. Software can check the spelling of all the words in a document faster than a person can find the first one by flipping through the pages of a printed dictionary. But a simple spell checker looks up each word only to discover whether it is in its dictionary. It does not check whether the writer uses the word properly.*

* Grammar checkers were rudimentary when the poem first circulated on the Internet. They would now catch some of the errors.



The convenience of using a computer system and abdication of responsibility to exercise judgment can encourage a mental laziness with serious consequences. A trucker in Britain got his truck stuck on a small farm road after ignoring a sign saying the road was not suitable for large vehicles. He was unquestioningly following the directions of a navigation system. A newspaper editor in Pakistan received a letter to the editor by email and inserted it into the newspaper without reading beyond the title. The letter was an attack on the prophet Muhammad. Angry Muslims set fires in the newspaper office. Several editors were arrested and charged with blasphemy, sometimes punishable by death.¹² Back when newspaper content was still typeset and copyedited, such an accident would have been unlikely.

Businesses make decisions about loan and insurance applications with the help of software that analyzes risks. School districts make decisions about the progress of students and the careers of administrators on the basis of computer-graded and -calibrated tests. They sometimes make bad decisions because of ignorance of the kinds of errors that limitations of the system can cause. Law enforcement agents arrested people when a check of an FBI database showed an arrest warrant for someone with a similar name. Do officers think that because the computer displayed the warrant, the system has decided that the person they are checking is the wanted person? Or does an officer understand that the system simply displays any close matches and that the responsibility for the arrest decision lies with the officer?

Sometimes reliance on a computer system rather than human judgment becomes “institutionalized” in the sense that an organization’s management and the legal system can exert strong pressure on individual professionals or employees to do what the computer says. In bureaucracies, a decision maker might feel that there is less personal risk (and less bother) in just accepting what the software produces rather than doing additional checking or making a decision the software does not support. Computer programs advise doctors on treatments for patients. It is critical to remember that, in complex fields, the computer systems might provide valuable information and ideas but might not be good enough to substitute for an experienced professional’s judgment. In some institutions, when something goes wrong, “I did what the program recommended” is a stronger defense (to superiors or against a lawsuit) than “I did what my professional judgment and

experience recommended.” Such institutions are encouraging abdication of personal responsibility, with potentially harmful results.

7.1.2 COMPUTER MODELS

Likeness to truth is not the same thing as truth.

—Peter L. Bernstein¹³

Evaluating models

Computer-generated predictions based on mathematical models of subjects with important social impact frequently appear in the news. Figure 7.1 shows a few examples. A mathematical model is a collection of data and equations describing, or simulating, characteristics and behavior of the thing studied. The models and simulations of interest to us here require so much data and/or computation that they must be run on computers. Researchers and engineers do extensive modeling to simulate both physical systems, such as the design for a new car or the flow of water in a river, and intangible systems, such as parts of the economy. Models allow us to simulate and investigate the possible effects of different designs, scenarios, and policies. They have obvious social and economic benefits: They help train operators of power plants, submarines, and airplanes. They enable us to consider alternatives and make better decisions, reducing waste, cost, and risk. They enable us to project trends and plan better for the future.

Although the models we consider are abstract (i.e., mathematical), the meaning of the word “model” here is similar to its meaning in “model airplane.” Models are simplifications. Model airplanes generally do not have an engine, and the wing flaps might not move. In a chemistry class, we could use sticks and balls to build models of molecules to help us understand their properties. The molecule models might not show the components of the individual atoms. Similarly, mathematical models do not include equations for every factor that could influence the outcome. They often include simplified

-
- Population growth
 - The cost of a proposed government program
 - The effects of second-hand smoke
 - When we will run out of a critical natural resource
 - The effects of a tax cut on the economy
 - The threat of global warming
 - When a big earthquake is likely to occur
-

Figure 7.1 Some problems studied with computer models.

equations because the correct ones are unknown or too complicated. For example, we use a constant known as the acceleration of gravity in a simple equation to determine when an object dropped from a high place will hit the ground. We ignore the effect of wind in the equation, but, on some days, wind could make a difference.

Physical models are usually not the same size as the real thing. Model planes are smaller; the molecule model is larger. In mathematical models, it is time rather than physical size that often differs from reality. Computations done on a computer to model a complex physical process in detail often take more time than the actual process takes. For models of long-range phenomena, such as population growth and climate change, the computation must take less time than the real phenomenon for the results to be useful.

Predictions from expensive computers and complex computer programs impress people, but models vary enormously in quality. Some are worthless. Others are very reliable. Politicians and special interest groups use model predictions to justify multibillion-dollar programs and laws with significant impact on the economy and the standard of living and choices of millions of people. It is important for both computer professionals and the general public to have some idea of what is in such computer programs, where their uncertainties and weaknesses might lie, and how to evaluate their claims. It is the professional and ethical responsibility of those who design and develop models for public issues to describe honestly and accurately the results, assumptions, and limitations of their models.

The following questions help us determine the accuracy and usefulness of a model.

1. How well do the modelers understand the underlying science or theory (be it physics, chemistry, economics, or whatever) of the system they are studying? How well understood are the relevant properties of the materials involved? How accurate and complete are the data?
2. Models necessarily involve assumptions and simplifications of reality. What are the assumptions and simplifications in the model?
3. How closely do the results or predictions of the model correspond with results from physical experiments or real experience?

Among three models developed to predict the change in health care costs that would result if the United States adopted a national health system, the predictions varied by hundreds of billions of dollars. Two of the models predicted large increases and one predicted a drastic decrease.¹⁴ Why was there such a difference? There are both political and technical reasons why models might not be accurate. Political reasons, especially for this example, are probably obvious. In addition to technical reasons that the questions above suggest (incomplete knowledge of the system being modeled, incomplete or inaccurate data, and faulty assumptions or oversimplification), other reasons are that computing power could be inadequate for the number of computations needed to model the full complexity of the system, and the difficulty, if not impossibility, of numerically quantifying variables that represent human values and choices.

-
- How many times do parents reuse a cloth diaper before discarding it? (Values ranged from 90 to 167.)
 - Should the model give credit for energy recovered from incineration of waste? Or does pollution from incineration counterbalance the benefit?
 - How many cloth diapers do parents use each time they change a baby? (Many parents use two at once for increased protection.) Numbers in the models ranged from 1.72 to 1.9.
 - How should the model count pesticides used in growing cotton?
-

Figure 7.2 Factors in diaper life cycle modeling.

Are reusable (washable cloth) diapers better for the environment than disposable diapers? When environmentalists proposed bans and taxes on disposable diapers, this controversy consumed almost as much energy as diaper manufacturing. Several modelers developed computer models to study the question. We call this particular kind of model a life cycle analysis. It attempts to consider the resource use and environmental effects of all aspects of the product, including manufacture, use, and disposal. To illustrate how difficult such a study might be, Figure 7.2 lists a few of the questions about which the modelers made assumptions. Depending on the assumptions, the conclusions differed.¹⁵ It is worth noting also that the models focused on one quality—environmental impact. To make a personal decision, we might consider the results of such a model (if we think it reliable), and we might also consider other factors such as cost, aesthetics, convenience, comfort, and health risks.

The U.S. Army Corps of Engineers uses mathematical models to predict how long an artificially constructed or replenished beach will last before waves wash it away. Two geologists have explained weaknesses in these models.¹⁶ Among other simplifying assumptions, the models assume that all waves have the same wavelength, that all waves come from the same direction, and that all grains of sand are the same size. A model uses only 6 of 49 parameters that might affect the amount of sand washed away. Even if these 6 are the most important (or if the model included all 49), the appropriate values for a particular beach are uncertain. Often, say the critics, the beaches do not last as long as the models predict, partly because the models do not accurately provide for relevant but irregular natural phenomena such as big storms.

Example: Modeling car crashes*

Car crash analysis programs use a technique called the finite-element method. They superimpose a grid on the frame of a car, dividing the car into a finite number of small

* An earlier version of this section appeared in my chapter, “Social and Legal Issues,” in *An Invitation to Computer Science* by G. Michael Schneider and Judith L. Gersting, West Publishing Co., 1995. (Used with permission.)

pieces, or elements. The grid is entered into the program, along with data describing the specifications of the materials making up each element (e.g., density, strength, and elasticity). Suppose we are studying the effects on the structure of the car from a head-on collision. Engineers initialize data to represent a crash into a wall at a specified speed. The program computes the force, acceleration, and displacement at each grid point and the stress and strain within each element. It repeats these calculations to show what happens as time passes in small increments. These programs require intensive computation to simulate 40–100 milliseconds of real time from the impact.

A real crash test can cost several hundred thousand dollars. It includes building and testing a unique prototype for a new car design. The crash analysis programs allow engineers to consider alternatives—for example, to vary the thickness of steel for selected components, or change materials altogether—and discover the effect without building another prototype for each alternative. But how good are the programs?

How well is the physics of car crashes understood? How accurate and complete are the data? Force and acceleration are basic principles. The physics involved in these programs is straightforward. Engineers know the relevant properties of steel, plastics, aluminum, glass, and other materials in a car fairly well. However, although they understand the behavior of the materials when force is applied gradually, they know less about the behavior of some materials under abrupt acceleration, as in a high-speed impact, and their behavior near or at breaking point. There are good data on the density, elasticity, and other characteristics of materials used in the model.

What simplifications do the programs make? The grid pattern is the most obvious. A car is smooth, not made up of little blocks. Also, time is continuous. It does not pass in discrete steps. The accuracy of a simulation depends in part on how fine the grid is and how small the time intervals are. Current computer speeds allow updating the calculations on fine grids with small time intervals (e.g., one millionth of a second).

How do the computed results compare to actual crash tests on real cars? High-speed cameras record real crash tests. Engineers attach sensors to the car and mark reference points on the frame. They compare the values the sensors record with values the program computes. They physically measure the distortion or displacement of the reference points, then compare these measurements to the computed positions of the points. Starting with the results of the physical crash, the engineers use elementary physics to calculate backward and determine the deceleration and other forces acting on the car. They compare these to the values computed in the simulation. The conclusion? Crash analysis programs do an extremely good job. In part because of the confidence that has developed over time in the validity of the results, engineers use variations of the same crash analysis modeling programs in a large variety of other impact applications, including those in Figure 7.3.

Engineers who work with the crash analysis programs do not believe that they will or should eliminate physical crash testing. The computer program is an implementation of theory. Results could be poor if something happens that the program designers simply did

-
- Predict damage to a hazardous waste container if dropped.
 - Predict damage to an airplane windshield or nacelle (engine covering) if hit by a bird.
 - Determine whether beer cans would get dented if an assembly line were speeded up.
 - Simulate a medical procedure called balloon angioplasty, where doctors insert a balloon in a blocked artery and inflate it to open the artery. The computer program helps researchers determine how to perform the procedure with less damage to the arterial wall.
 - Predict the action of airbags and the proper location for sensors that inflate them.
 - Design interior parts of cars to reduce injuries during crashes (e.g., from the impact of a steering wheel on a human chest).
 - Design bicycle and motorcycle helmets to reduce head injuries.
 - Design cameras to reduce damage if dropped.
 - Forecast effects of earthquakes on bridges and buildings.
-

Figure 7.3 Other uses of crash analysis programs.

not consider. The crash analysis programs are excellent design tools that enable increases in safety with far less development cost. The physical crash test is confirmation.

Example: Modeling climate

The earth has Ice Ages and warm interglacial periods. We are now in an interglacial period that is more than 11,000 years old. Within such periods, climate varies over time and in different parts of the world. For example, the Northern Hemisphere experienced both a medieval warm period about a thousand years ago and a later colder period (roughly 1550–1850), sometimes called the Little Ice Age.

Climate change is a very complex phenomenon. Solar radiation warms the earth. Some of the heat is reflected back, and gases trap some in the atmosphere. The latter phenomenon is known as the greenhouse effect. Without it, the temperature on the earth would be too cold to support life. Water vapor is the main greenhouse gas, but there are several other significant greenhouse gases as well. Among those whose concentration has been increased by human activity (in particular, burning of fossil fuels), carbon dioxide (CO₂) is most important. An upward trend in CO₂ concentration began roughly 16,000 years ago. However, since the beginning of the Industrial Revolution, CO₂ concentration has been increasing at a faster rate.* Between the period 1850–1899 and the period 2001–

* CO₂ concentration has increased by almost 40% since 1750.¹⁷ The older data come from measurements of gases trapped in ice cores drilled in Antarctica and Greenland.

2005, average global temperature rose roughly 0.76°C .^{*18} The increase has been steeper since roughly 1980. The global temperature increase raised concern about the threat of excess global warming, possibly caused by human-induced increase of CO_2 and other greenhouse gases in the atmosphere. Global warming predictions are based on computer models of climate. We consider those models. Since 1990, the Intergovernmental Panel on Climate Change (IPCC), sponsored by the United Nations and the World Meteorological Organization, has published comprehensive reports on the science of climate change and the quality and projections of climate models roughly every five years. Much of the information in this section comes from those reports.¹⁹

Climate models, like the car crash analysis models, calculate relevant variables for grid points and elements (grid boxes) for specified simulated time intervals. The grid circles the earth, rises through the atmosphere, and goes down into the ocean. The models contain information about the sun's energy output; the orbit, inclination, and rotation of the earth; geography (a map of land masses); topography (mountains, etc.); clouds; sea and polar ice; soil and air moisture; and a large number of other factors. Equations simulate atmospheric pressure, temperature, wind speed and direction, moisture, precipitation, ocean currents, and so forth. Researchers use climate models to study several aspects of future climate. They try, for example, to determine the effect of doubling CO_2 concentration in the atmosphere. (Current trends suggest the concentration will have doubled, from its approximate level at the beginning of the 20th century, by some time in the 21st century.[†]) Models also project the likely increase in global temperature, sea level, and other climate characteristics in various scenarios with assumptions about population, industrial and economic activity, energy use, and so on, for the rest of this century. Another task for the models is to distinguish how much warming is caused by human activity and how much is natural. The IPCC has concluded that it is "extremely likely" that human activity has had a substantial warming effect on climate since 1750.²⁰

Climate models have improved over the few decades that scientists have been developing and working with them. The models used in the 1980s and 1990s were quite limited. Here is a brief sampling of simplifications, assumptions, and factors modelers did not fully understand: The models did not distinguish day and night.²¹ They used a fairly coarse grid (with points roughly 500 kilometers apart). They did not include the El Niño phenomenon. They made assumptions about methane (a greenhouse gas) that scientists later determined were incorrect. They did not include aerosols (small particles in the air) that have a cooling effect. Clouds are extremely important to climate, but many processes involved with the formation, effects, and dissipation of clouds were not particularly well understood. The IPCC summarized in 2001: "As has been the case since the first IPCC Assessment Report in 1990, probably the greatest uncertainty in future projections of

* With an error range of ± 0.19 .

† Other greenhouse gases are included too by converting their amount and effect to the equivalent number of units of CO_2 .

climate arises from clouds and their interactions with radiation. . . . Clouds represent a significant source of potential error in climate simulations.”²² The extremely simplified representations of the oceans in these models was another very significant weakness. Computing power was insufficient to do the many calculations to simulate ocean behavior. When run on past data, some of the early climate models predicted temperature increases three to five times as high as what actually occurred over the previous century. The 1990 IPCC report predicted that temperature would increase 0.3°C per decade (with an error range of 0.2°–0.5°C). The actual temperature increase over the next two decades was lower than that.²³ Thus, it should not be surprising that there was much skepticism about the climate models and their projections.

Current models are more detailed and complex. Increased computer power allows the use of finer grids (with points spaced roughly 100–300 kilometers apart) and more experiments with the models. Increased data collection and basic science research have been improving the understanding of the behavior and interactions of climate system components. The models project that doubling the concentration of greenhouse gases in the atmosphere will cause a global temperature increase within the range 2°–4.5°C. The models project warming of 0.2°C per decade for the next few decades and a sea level rise of between 8 and 23 inches by the end of the 21st century.²⁴

How well is the science understood? How accurate are the data? Climatologists know an enormous amount about climate. The models incorporate a huge amount of good science and data. But the amount not known is also large.

Much of the variation in model results comes from the still troublesome lack of full understanding of clouds. When the earth warms, water evaporates, and the additional water vapor in the atmosphere absorbs more thermal energy, warming the atmosphere farther. On the other hand, water vapor forms clouds, which reflect incoming solar radiation with a cooling effect. Thus, clouds have positive (destabilizing) and negative (stabilizing) feedback effects. The basic science of the mechanisms is fairly well understood, but not the complexity and magnitude of the feedbacks.²⁵

A related area of uncertainty has been the impact of variations in output from the sun. Recent research on the interactions between solar activity, cosmic rays (radiation from space), and cloud formation suggests that solar activity might have an impact on cloud formation, and thus on warming, that the climate models do not include. The research is at an early stage, the theory is controversial, and the magnitude of the impact is unknown. Experiments are continuing.²⁶

There is insufficient data on many phenomena for the period before satellites collected data. For example, the IPCC lists among “key uncertainties” insufficient data to draw conclusions about trends in thickness of Antarctic sea ice.

The temperature data sets that models use for temperature over the past century have been a source of some contention. They include many kinds of limitations (for example, few monitoring stations in the oceans and remote land areas) and errors. In 2011, the Berkeley Earth Surface Temperature project completed a multiyear effort analyzing

temperature data from 15 different data sets. It reviewed algorithms and statistical methods used to develop the data sets, and it developed new statistical methods to try to overcome problems in the previous methods. It developed and published a new global surface temperature record and provided an uncertainty analysis.²⁷ Climate science researchers had much praise for the quality of the work in this project and for Berkeley Earth's decision to openly publish its methodology.

What are the assumptions and simplifications in the models? Ideally, equations derived from the underlying science (generally, physics and chemistry) would model all the processes that affect climate. This is not possible, because it would require too much computation time and because all the underlying science is not known. Simplified equations, called parametrizations, represent many processes; they seem to give realistic results but are not derived from scientific theory. The specific parametrizations vary among the models, reflecting the choices of the modelers.

The IPCC acknowledges that the underlying complexity of the problem still hampers the accuracy of projections for future climate change. That is, even the extremely powerful computers of today are not sufficient to achieve an ideal level of resolution (grid size) and to include simulation of more processes that affect climate.²⁸

The model projections based on scenarios (rather than a specific increase in greenhouse gas concentration) include numerous assumptions about technological development, political control of emissions, population, economic development, energy use, and so on, throughout a century.

Science and fiction

Why do science fiction movies about global warming show the buildings of cities underwater? The entertainment industry exaggerates and dramatizes, of course. Why does an exhibit in a science museum show water up to the middle of the Statue of Liberty (about 200 feet above sea level)? A climate scientist once said: “[T]o capture the public’s imagination,” “we have to offer up scary scenarios, make simplified dramatic statements, and make little mention of any doubts we may have. . . . Each of us has to decide what the right balance is between being effective and being honest.”²⁹ Although he said he hoped climate scientists

could be both effective and honest, there is clearly an ethical problem when we trade honesty for something else. Is it a good idea? A 20-inch rise in sea level would be a very significant problem, but one we can tackle. Tens or hundreds of feet of sea level rise would be an enormous disaster. Exaggeration might lead people to take constructive action. Or exaggeration might lead to overreaction and counterproductive, expensive actions, draining resources from effective approaches. If we hope to solve real potential problems (such as flooding in low-lying areas), we must first identify them accurately.

How well do predictions of the models correspond with actual experience? The models predict seasonal variations and other actual broad-scale phenomena. The general patterns of predictions by different models are similar. For example, they all predict warming, and they all predict that more of the warming would take place near the poles and in winter. Many models now do a good job predicting air temperature near the surface of the earth (that is, close to observed temperatures) for the recent past. The models do well enough that the IPCC expresses many of its projections as very likely or likely.

For more than a decade at the beginning of the 21st century, global temperature fluctuated but did not rise overall. The models did not indicate that this would happen. Scientists are devising and testing theories to account for it. The amount of water vapor in the stratosphere is a key suspect.³⁰ Models designed to project long-term trends might not predict short-term variations well. Thus, we do not know yet if the first decade of the century was a short-term variation or whether it will require revisions in the models.

7.2 The “Digital Divide”³¹

The term *digital divide* refers to the fact that some groups of people (the “haves”) enjoy access to and regularly use the various forms of modern information technology, while others (the “have-nots”) do not. The focus of the discussion about “the digital divide” has shifted over time. In the 1990s, the focus was on access to computers and the Internet for poor people, people in rural areas, and certain demographic groups within the United States (and other developed countries). As more people acquired digital technology and Internet access, focus shifted to a divide among those who have broadband and those who do not. There is also more focus now on the digital divide between developed countries and poor countries.

7.2.1 TRENDS IN COMPUTER ACCESS

Once upon a time, everyone in the world had equal access to personal computers and the Internet. They did not exist, and we all had none. Later, a small, elite minority enjoyed these new, expensive tools. As the technology began to spread and its value became clearer, people became more concerned about the gap in access. Poor children and children of some ethnic minorities had less access to computers both in schools and at home. In the early 1990s, only about 10% of Net users were women. By 1997, the gender gap had vanished,³² but other gaps remained.* Black and Hispanic households were about half as likely as the general population to own a computer. Access in rural and remote regions lagged behind the cities.

* A gender gap remains among those who work in information technology fields. Only about one-quarter are women, and the percentage of women undergraduates interested in the computer science major dropped drastically between 2000 and 2009.

Cost is one factor that affects access by the general population. Ease of use is another. At first, personal computers and the Internet were difficult to use. Software innovations, such as point-and-click graphical user interfaces, Web browsers, and search engines made computer use significantly more comfortable for ordinary people. With lower prices, more useful applications, and ease of use, ownership and access spread quickly. The data I found about the extent of computer ownership and Web access differ in specific numbers, but all showed the same trends: In 1990, 22% of households in the United States owned a computer. In 2001, 84% of homes with children in middle and high school had Internet access. That was a significantly higher percentage of households overall, suggesting that families perceived access to be important for their children and allocated their spending accordingly.³³

Individuals, businesses, community organizations, foundations, and government programs contributed to the spread of computers and Internet access. Internet cafés sprang up from Alaska to Cairo in the 1990s when Net access from home was relatively uncommon. The federal government and local governments spent billions of dollars on technology for schools. By the end of the 1990s, most public libraries provided Internet access for the public for free. By 2000, 98% of high schools had Internet access. At about the same time, African Americans, people 65 and older, and Hispanics increased their use of the Internet significantly. Groups with low access in earlier years began to catch up. The gaps among Hispanic, black, and white people almost completely disappeared among those with the same education levels.³⁴ By 2011, there were more than 300 million cellphone subscriptions in the United States. Computing and new communication technology reached more households much more quickly than earlier technologies such as telephone, television, electricity, and automobiles.

Virtually all technological innovation is first available to the rich (or others willing to pay the initially high price). The early purchases finance improvements in design and production techniques that bring the price down. Prices of many consumer products follow this pattern. Telephones and televisions were originally luxuries of the rich. Now, almost everyone in developed countries has them. (By 2006, there were more televisions than people in the United States.) When first introduced in the 1980s, compact disk music players cost \$1000. Now we play music and video on our phones. Computer prices plunged more dramatically than prices of most other products, even while the memory, speed, and variety of input/output devices and software increased enormously.* The vast resources of the Internet are available for about what home telephone service used to cost. The phenomenon that new technologies and inventions first are expensive luxuries, then become cheaper and spread throughout the population, has led some observers to conclude that it is more accurate to think of people as “haves” and “have-laters” rather than “haves” and “have-nots.”³⁵

* For example, the cost of disk storage fell from hundreds of dollars per megabyte in the 1980s to about \$100 a terabyte (one million times as much space) by 2012.

Access to broadband connection is a newer version of a digital divide. The same demographic groups that were the have-laters with respect to Internet access in the 1990s had less broadband access in the first decade of the 2000s, with similar disadvantages. Without broadband access, it is more difficult to find employment opportunities, access news and information, and make use of some online health information and tools. Children are less likely to graduate from high school. People without broadband are less likely to create Internet content. According to Connected Nation, in 2011 only 46% of low-income households with children and 37% of low-income minority households with children had broadband at home, compared to 66% of households nationally. (About 20% of all U.S. households had broadband in 2003.) The largest barrier to adopting broadband, according to people surveyed, was cost, followed by digital literacy and not believing that the Internet was relevant to their lives. In 2011, the FCC approved a program to extend broadband service to rural areas of the United States. Businesses and nonprofit organizations started a program to address cost and digital literacy. Under this initiative, major cable carriers offer broadband service at a low price to low-income customers. Best Buy, Microsoft, America’s Promise Alliance, and United Way, among many others, provide digital literacy training.³⁶

A related digital divide exists among content consumers and content producers on the Internet. Internet users create blogs, Web pages, videos, and product reviews. Being a content creator empowers a user to communicate his or her message to a large number of people. The Internet can be a strong agent for change for those who have the skills, education, and tools to create content. Content creators tend to be people who access the Internet frequently from multiple places using multiple gadgets. They also are more educated.³⁷ The content-production divide shows a gap among users based on socioeconomic status. How should we view the Internet content-production divide? Before the Internet, a very small percent of people wrote books and articles and produced movies and television shows. The vast majority of people were content consumers only. Is the current divide less of a social problem than the pre-Internet divide because so many more people can now create content, or is it more of a problem than before because it isolates a smaller part of the population that cannot?

7.2.2 THE GLOBAL DIVIDE AND THE NEXT BILLION USERS

Approximately two billion people worldwide use the Internet, a fivefold increase over roughly a decade.³⁸ From one perspective, that is an extraordinary accomplishment in a very short time. From another perspective, it means that about five billion people do not use the Internet. Lack of access to the Internet in much of the world has the same causes as lack of health care, education, and so on: poverty, isolation, poor economies, and politics.

Both nonprofit organizations and huge computer companies have ongoing projects to spread computer access to more people in developing countries. Some companies use

the catchphrase “the next billion users” to describe the people their programs address. For the companies, these programs create good will and—if successful in improving the standard of living and economies of the target countries—a large future customer base. Companies have trained hundreds of thousands of teachers to use technology effectively in classrooms in China, India, and other countries.

One Laptop per Child is a nonprofit organization that supplies an inexpensive laptop computer specially designed for elementary school children in developing countries. The laptop works in extreme heat or cold, extremes of humidity, and dusty or rainy environments. The power requirements are very low. The program provided an important lesson: giving out computers and walking away will not close the digital divide. The success of the program in implementing the technology into school curricula depends on the presence of supporting social and technical infrastructures such as electricity, networks, tech support, parental support, teacher attitudes towards technology, and administrative school support. Purely financial resources can be less important than these factors.³⁹

Some people active in movements to shrink the “digital divide” emphasize the need to provide access in ways appropriate to the local culture. For example, one website argues that access can hurt the poor “by loosening the bonds of tradition.” In many countries, access “is one-way, entertainment-oriented, commercial.” Access might accelerate the exodus of untrained, unprepared young people from rural areas into cities.⁴⁰ How significant are these concerns? What can be done to alleviate them?

Only a few years ago, most people in the world had never made a telephone call. By the end of 2010, there were five billion cellphone subscriptions.*⁴¹ Almost every time I have read about a program to bring Internet access or cellphones to rural, third-world adults over the past several years, the most immediate uses are similar. Farmers use the Internet to learn about better farming techniques and to get up-to-date pricing information for their crops. Fishermen use their cellphones to find a nearby village where they will get a good price for their catch. As the technology spreads, food production and economic well-being improve. Some see each new digital divide as a serious social problem. What is perhaps most surprising is how quickly most of these divides shrink and how much more quickly they shrink than did previous technological divides between rich and poor, men and women, black and white, or developed and undeveloped countries.

7.3 Neo-Luddite Views of Computers, Technology, and Quality of Life

The microchip is . . . made of silicon, or sand—a natural resource that is in great abundance and has virtually no monetary value. Yet the

* That does not mean that five billion people had cellphones. In wealthier areas, some people have more than one.

combination of a few grains of this sand and the infinite inventiveness of the human mind has led to the creation of a machine that will both create trillions of dollars of added wealth for the inhabitants of the earth in the next century and will do so with incomprehensibly vast savings in physical labor and natural resources.

—Stephen Moore⁴²

Quite apart from the environmental and medical evils associated with them being produced and used, there are two moral judgments against computers. One is that computerization enables the large forces of our civilization to operate more swiftly and efficiently in their pernicious goals of making money and producing things . . . And secondly, in the course of using these, these forces are destroying nature with more speed and efficiency than ever before.

—Kirkpatrick Sale⁴³

7.3.1 CRITICISMS OF COMPUTING TECHNOLOGIES

The quotations above, both from 1995, illustrate the extreme divergence of views about the anticipated value of computer technology. Evaluations cover the spectrum from “miracle” to “catastrophe.” Although most of this book discusses problems that arise with the use of computers, the Internet, and other digital communications technologies, the implicit (and sometimes explicit) view has been that these technologies are a positive development bringing us many benefits. The potential for loss of freedom and privacy via government surveillance and the building of consumer dossiers is a serious danger. Computer crime is expensive, and changes in employment are disruptive. Our discussion of systems failures in the next chapter warns us that some potential applications can have horrifying risks. We might urgently try to prevent implementation of some applications and urgently advocate for increased protection from risks, yet not consider the threats and risks as reasons for condemning the technology as a whole. For the most part, we have looked at new risks and negative side effects as problems that occur in the natural process of change, either problems we need to solve or the price we pay for the benefits, part of a trade-off. Many people with quite different political views share this attitude, although they disagree about the significance of specific computer-related problems and about exactly how to solve them.

On the other hand, there are people who utterly reject the view that computing technology is a positive development with many important benefits. They see the benefits as few and overwhelmingly outweighed by the damage done. Neil Postman says that

voting, shopping, banking, and getting information at home online is a “catastrophe.” There are fewer opportunities for people to be “co-present,” resulting in isolation from neighbors. Richard Sclove and Jeffrey Scheuer argue that electronic communication will erode family and community life to the point that people will mourn the loss of depth and meaning in their lives.⁴⁴ A comment made by one reviewer of this book illustrates the difference in perspective. He objected to the “gift of fire” analogy I use to suggest that computers can be very useful and also very dangerous. The reviewer thought “Pandora’s box” was more appropriate. Pandora’s box held “all the ills of mankind.” Kirkpatrick Sale, author of *Rebels Against the Future*, used to demonstrate his opinion of computers by smashing one with a sledgehammer at public appearances.

In England in 1811–1812, people burned factories and mills in efforts to stop the technologies and social changes that were eliminating their jobs. Many were weavers who had worked at home on small machines. They were called Luddites.* For 200 years, the memory of the violent Luddite uprising has endured as the most dramatic symbol of opposition to the Industrial Revolution. The term “Luddite” has long been a derisive description for people who oppose technological progress. More recently, critics of technology have adopted it as an honorable term. Kirkpatrick Sale and many others who share his viewpoint call themselves neo-Luddites, or simply Luddites.

What do the neo-Luddites find so reprehensible about computers? Some of their criticisms are problems that also trouble people whose view of computing technology is generally positive, problems we discussed in earlier chapters. One of the differentiating characteristics of the neo-Luddites is that they focus on these problems, seeing no solutions or trade-offs, and conclude that computers are a terribly bad development for humankind. Among their specific criticisms are the following:

- Computers cause massive unemployment and de-skilling of jobs. “Sweatshop labor is involved in their manufacture.”⁴⁵
- Computers “manufacture needs”; that is, we use them just because they are there, not because they satisfy real needs.
- Computers cause social inequity.
- Computers cause social disintegration; they are dehumanizing. They weaken communities and lead to isolation of people from each other.
- Computers separate humans from nature and destroy the environment.
- Computers benefit big business and big government most.
- Use of computers in schools thwarts development of social skills, human values, and intellectual skills in children. They create an “ominous uniformity of knowledge” consistent with corporate values.⁴⁶

* The name Luddite comes from General Ned Ludd, the fictitious, symbolic leader of the movement.

- Computers do little or nothing to solve real human problems. For example, Neil Postman, in response to claims of the benefits of access to information, argues that “if families break up, children are mistreated, crime terrorizes a city, education is impotent, it does not happen because of inadequate information.”⁴⁷

Some of these criticisms might seem unfair. The conditions in computer factories hardly compare to conditions in the sweatshop factories of the early Industrial Revolution. In Chapter 6, we saw that computers eliminate some jobs, and that the pace of computerization causes disruptions, but the case that computers, and technology in general, cause massive unemployment is not convincing. Blaming computers for social inequity in the world ignores thousands of years of history. Postman is right that inadequate information is not the source of most social problems. A computer in the classroom does not replace good parents in the home. But should this be a criticism of computers and information systems? Access to information and communication can assist in solving problems and is not likely to hurt. The main problem for ordinary people, Postman says, is how to find meaning in life. We need answers to questions like “Why are we here?” and “How are we supposed to behave?”⁴⁸ Is it a valid criticism of computing technology that it does not solve fundamental social and philosophical problems that have engaged us for thousands of years?

To the neo-Luddites, the view that computers are fundamentally malevolent is part of a wider view that almost all of technology is malevolent. To the modern-day Luddites, computer technology is just the latest, but in many ways the worst, stage in the decline of what was good in human society. Computers are worse than earlier technologies because of their enormous speed and flexibility. Computers increase the negative trends that technology causes. Thus, if one points out that a particular problem blamed on computers already existed because of an earlier technology, Luddites consider the distinction to be a minor one.

The depth of the antipathy to technology in the Luddite view is perhaps made clearer by attitudes toward common devices most of us use daily. For example, Sale has said, “I find talking on the phone a physical pain, as well as a mental anguish.” Sven Birkerts, another critic of computers, says that if he lived in 1900, he would probably have opposed the telephone.* Speaking of the invention of the printing press, Sale laments that “literacy . . . destroys orality.” He regards not only computers but civilization as a catastrophe. Some of us see modern medicine as a life-saving and life-enhancing boon to humanity; some Luddites point out that it gave us the population explosion and extended senility.⁴⁹

* Critics of telephones complained that they replaced true human interaction with disembodied, remote voices. They actually expanded and deepened social relationships for isolated people—for example, women in general (farm wives, in particular) and the elderly.⁵⁰

Having read and listened to the arguments of technology enthusiasts and technology critics, I find it striking that different people look at the same history, the same society, the same products and services, the same jobs—and come to diametrically opposed conclusions about what they see. There is a fundamental difference between the world views of supporters and opponents of technology. It is more than the difference between seeing a glass as half full or half empty. The difference seems to be one of contrasting views about what should be in the glass. Supporters of technology see an upward trend in quality of life, beginning with people living at the mercy of nature with an empty glass that technology has been gradually filling. Neo-Luddites view the glass as originally full when people lived in small communities with little impact on nature; they see technology as draining the glass.

The neo-Luddite view is associated with a particular view of the appropriate way of life for human beings. For example, Sale's first point, in the quotation at the beginning of this section, makes the moral judgment that making money and producing things is pernicious. His introductory remark and his second point barely hint at the unusually high valuation he places on not disturbing nature (unusually high even in the contemporary context, where there is much awareness of the importance of protecting the environment). We explore these views further.

7.3.2 VIEWS OF ECONOMICS, NATURE, AND HUMAN NEEDS

Luddites generally have a negative view of capitalism, business, markets, consumer products, factories, and modern forms of work. They see the profit-seeking goals of businesses as in fundamental conflict with the well-being of workers and the natural environment. They see work in factories, large offices, and business in general as dehumanizing, dreary, and bad for the health of the workers. Hence, for example, the Luddite criticisms of the clock. Neil Postman describes the invention of the clock as “the technology of greatest use to men who wished to devote themselves to the accumulation of money.”⁵¹

Choice of words, making subtle differences in a statement, sometimes illustrate the difference in perspective between Luddites and non-Luddites. What is the purpose of technology? To the Luddites, it is to eliminate jobs to reduce the costs of production. To proponents of technology, it is to reduce the effort needed to produce goods and services. The two statements say nearly the same thing, but the first suggests massive unemployment, profits for capitalists, and a poorer life for most workers. The second suggests improvements in wealth and the standard of living.

The Luddite view combines a negative attitude toward business with a high estimation of the power of corporations to manipulate and control workers and consumers. For example, Richard Sclove describes telecommuting as being “imposed by business.” (Interestingly, one of the common criticisms of the Industrial Revolution was that working in factories instead of at home weakened families and local community.)

Luddites make particularly strong criticisms of automobiles, of cities, and of most technologies involved in communications and transportation. Thus, it is worth noting that most of us get both personal and social benefits from them. Cities are centers of culture, wealth production, education, and job opportunities.⁵² Modern transportation and communication reduce the price of products and increase their variety and availability. For example, we can eat fresh fruits and vegetables all year. We can look up menus and movie schedules on our smartphone to find what we want. We can shop worldwide on the Web. We can commute a long distance to take a better job without having to sell our house and move. If we move to a new city for college or a job, modern conveniences such as airplanes, telephones, and the Internet make the separations less unpleasant. We can visit more often in person. We can share greetings and activities with friends and family members via social media. Luddites and other critics of technology do not value these advantages highly. In some cases, in their point of view, the advantages are merely ameliorating other problems technology causes. For example, Postman quotes Sigmund Freud's comment, "If there had been no railway to conquer distances, my child would never have left his native town and I should need no telephone to hear his voice."⁵³

Does the technology create the need for itself?

A common criticism of capitalism is that it survives by convincing us to buy products we do not need. Luddites argue, similarly, that technology causes production of things we do not need. This contrasts with the market-oriented view that sees consumer choices as determining which products, services, and businesses succeed or fail (in the absence of government favoritism, subsidies, and restrictions). We examine the issue of created needs.

Sale argued that small, portable computers do not "meet any known or expressed need," but companies produced them simply because miniaturization of computing components made it possible. People have bought many millions of laptops, tablet computers, and cellphones. The number of uses is phenomenal. So, does a mobile computer meet a need? It depends on what we mean by "need." Do we need to do homework in the backyard or listen to music on an iPod? Does an architect or contractor need a laptop at a construction site? Those who emphasize the value of individual action and choices argue that needs are relative to goals, and goals are held by individuals. Thus, should we ask whether "we," as a society, need portable computers? Or should this be an individual decision with different responses? Many people demonstrate, by their purchases, that they want portable computers. Anyone who does not feel a desire or need for one does not have to buy one. The Luddites, who believe that advertising, work pressure, or other external forces manipulate buyers, reject this individual-oriented approach.

Environmental and anti-technology groups use computers and the Web. The editor of *Wild Earth*, who considers himself a neo-Luddite, said he "inclines toward the view

Wal-Mart and e-commerce versus downtown and community⁵⁴

Does electronic commerce force changes on communities that no one wants? Richard Sclove and Jeffrey Scheuer think so.⁵⁵ They use the analogy of a Wal-Mart store draining business from downtown shops, resulting in the decline of the downtown community, a “result that no consumers wanted or intended.” They generalize from the Wal-Mart scenario to cyberspace. As we conduct more economic transactions electronically, we lose more local stores, local professional and social services, and convivial public spaces like the downtowns of small towns. Consumers are “compelled” to use electronic services, “like it or not.” Other strong critics of technology share the underlying point of view of Sclove and Scheuer, so it is worth examining their argument.

The Wal-Mart analogy is a good one. The scenario is useful for illustrating and clarifying some issues about the impact of e-commerce on communities. Suppose, say Sclove and Scheuer, that a new Wal-Mart store has opened just outside of town and about half the town residents begin to do about a third of their shopping there, while the others continue to do all their shopping downtown. Everyone shops downtown, and everyone wants the downtown stores to remain. But downtown stores have lost about 16.5% of their sales, and many will not survive. Sclove and Scheuer describe this as an “involuntary transformation” that no consumer wanted or intended. It occurs, they say, because of a “perverse market dynamic.” The changes, however, are not involuntary or perverse. The core of the problem with Sclove’s and Scheuer’s interpretation is their failure to make two important distinctions: the distinction between wanting something and the willingness to pay for it, and the distinction between something being coerced or involuntary, on the one hand, and being unwanted, unintended, or unexpected on the other.

Consider a simpler situation for a moment. Suppose we poll the adult residents of a small town with a population of, say, 3000 and ask if they would like to have a fine French restaurant in town. Almost everyone says yes. Will a French restaurant open in the town? Probably not. Almost everyone wants it, yet there is not enough potential business for it to survive. There is a market dynamic at work, but it is not perverse. The fact that consumers want a particular service, store, or product is irrelevant if not enough people are willing to pay the prices that make the business viable. In Sclove’s

and Scheuer’s Wal-Mart scenario, the downtown stores could stay in business if the people were willing to pay higher prices to make up for the 16.5% of revenue lost to Wal-Mart. But we know that if the stores raise prices, they will almost certainly lose even more customers. The town residents are not willing to pay what it costs to keep the downtown stores in business. You might object: The townspeople did not have to pay the higher prices before. Why now? Because now the people who shop at Wal-Mart—or online—*have another choice*. Whatever price advantage or convenience lured them, they were not getting that benefit before. Again, a market dynamic is at work, but not a perverse one: competition.

The second issue about the Wal-Mart/e-commerce scenario is whether the change is an “involuntary” transformation. Sclove and Scheuer say that, as local businesses decline, people will be compelled to use electronic services, like it or not. Is this accurate? No more so than Wal-Mart shoppers or cyberspace enthusiasts were compelled to shop downtown (or from other offline stores), like it or not, before they had the new option. The new status quo is no more involuntary than the previous one. Although no one wants to see the downtown decline, the actions that could lead to that result are all voluntary. When a new store opens (online or offline), no one is forced to shop there. The impact on the downtown stores might not have been obvious to all the townspeople at the beginning (although now it is common enough that they might anticipate it), but an unexpected or unintended result is not the same as a coerced result. In a free society, individuals make millions of decisions based on their knowledge and preferences. This decentralized, individualized decision making produces a constantly changing pattern of stores, services, and investments (not to mention social and cultural patterns). No one can predict exactly what the result will be, and no one intends a particular picture of the economy or society, but (apart from government subsidies, prohibitions, and regulations) the actions of the consumers and merchants are voluntary. No one person can expect to have exactly the mix of shopping options (or other community characteristics) that he or she wants. If the result flows from the myriad decisions that consumers and producers make, it is not coerced. It is the process, not the result, that tells us whether an outside force is coercing people.

Do we need cellphones?

Hundreds of thousands of people have heart attacks in the United States each year. Treatment received in the first few minutes can be critical to their survival. A fire department in California helped develop a smartphone app that alerts people trained in CPR if they are near the location where a person is having a

heart attack, perhaps in the same office building, shopping center, or neighborhood. The app provides the location of the victim and the locations of any nearby emergency defibrillator devices, so a trained person can get to the scene quickly and has the tools to save a life.

that technology is inherently evil,” but he “disseminates this view via E-mail, computer, and laser printer.”⁵⁶ An interviewer reported that in 2007, after a long career attacking computers, Kirkpatrick Sale was using a laptop. The question is: Are Sale and the editor of *Wild Earth* using computer equipment because of an artificial need or because it is useful and helpful to them? Sale sees the use of computers as an uncomfortable compromise. The use of computers, he says, insidiously embeds into the user the values and thought processes of the society that makes the technology.⁵⁷

The argument that capitalists or technologies manipulate people to buy things they do not really want, like the argument that use of computers has an insidiously corrupting effect on computer users, displays a low view of the judgment and autonomy of ordinary people. It is one thing to differ with another person's values and choices. It is another to conclude that, because of the difference, the other person is weak and incapable of making his or her own decisions. The Luddite view of the appropriate way of life puts little value on modern comforts and conveniences or on the availability of a large variety of goods and services. Perhaps most people value these things more highly than the Luddites do. To get a clearer understanding of the Luddite view of a proper life style, we consider some of their comments on the relationship of humans and nature.

Nature and human life styles

Luddites argue that technology has made no improvement in life, or at best improvements of little importance. Sale's list of benefits includes speed, ease, and mass access—all of which he disdains. Sale says that although individuals might feel their lives are better because of computers, the perceived benefits are “industrial virtues that may not be virtues in another morality.” He defines moral judgment as “the capacity to decide that a thing is right when it enhances the integrity, stability, and beauty of nature and is wrong when it does otherwise.”⁵⁸ Jerry Mander, founder of the Center for Deep Ecology and author of books critical of technology and globalization, points out that thousands of generations of humans got along without computers, suggesting that we could do just fine without them too. While some people evaluate trade-offs between negative side effects of pesticides

and the benefits of reducing diseases or protecting food crops, Mander's objections to technology lead him to the conclusion that there can be no "good" pesticide. While many people work on technological, legal, and educational approaches to reducing the gasoline usage of automobiles, Mander says there can be no "good" automobile.⁵⁹

What are the underlying premises behind these comments by Sale and Mander? We consider Sale's comment on moral judgment first. Many debates about the environment set up a humans-versus-nature dichotomy.⁶⁰ This is not the true conflict. Nature, biodiversity, forests, a hospitable climate, clean air and water, open space away from cities—these are all important and valuable to humans. So is shelter from the rain, cold, and heat. So are life-saving medicines and medical techniques. Conflicts about the environment are not conflicts between humans and nature. They are conflicts between people with different views about how to meet human needs. In contrast to Sale's statement, moral judgment, to many people, and for many centuries, has meant the capacity to choose that which enhances human life, reduces misery, and increases freedom and happiness. Sale's comment chooses nature, not humanity, as the primary standard of moral value.

Whether an automobile (or computing device) is "good," by a human-centered standard, depends on whether it meets our needs, how well it does so, at what cost (to the environment and society, as well as to our bank account), and how well it compares to alternatives. Critics of modern technologies point out their weaknesses but often ignore the weaknesses of alternatives—for example, the millions of acres once needed to grow feed for horses and the hundreds of tons of horse manure dropped on the streets of cities each day, a century ago.⁶¹ Mander's comment about automobiles again raises the issues of our standard of value and our need for a product or service. Candles, gas lamps, and kerosene lamps filled homes with fumes and soot. Do we need electricity? Do we need hot water on tap, movies, and symphony orchestras? Or do we need nothing more than food and shelter? Do we need an average life expectancy of more than 25 years? Do we want to merely exist—do we *need* even that?—or do we want long, happy, comfortable lives filled with time for love, interesting activities, and an opportunity to use our marvelously inventive brains?

*The Web is alive, and filled with life, nearly as complex and, well,
natural as a primordial swamp.*

—John Perry Barlow⁶²

Accomplishments of technology

It is easy to miss the extreme changes in quality of life that have taken place over the past few centuries. We mention here a scattering of examples.

Technology and the Industrial Revolution have had a dramatic impact on life expectancy. A study in 1662 estimated that only 25% of people in London lived to age

Environmental impacts of computing technology

I had thought of including a section in this book on environmental impacts of computers, mobile devices, and the Internet. As I looked for data, I concluded that attempts to quantify environmental benefits and costs would be subject to the same weaknesses and criticisms of models that we discussed in Section 7.1.2. It is extremely difficult to measure impacts and to determine how to compare to impacts of technologies and activities that computing technology replaces. However, we can make some observations.

Production of computers is energy intensive and uses hazardous materials. Because of these materials, disposal is an issue, as it is for fluorescent light bulbs. Running and cooling the millions of servers on the Internet in the United States accounts for about 2% of U.S. electric power usage,⁶³ more than the U.S. auto industry and less than the chemical industry. There are estimates that production of computers uses roughly twice as much energy as operating them.

On the other hand, digitally controlled machinery uses less power than older electromechanical controls. Digital sensors and controls for regulating lighting, heating, air conditioning, and farm irrigation (among many other examples) save resources by determining just what is needed and thus reduce waste. Microprocessors control hybrid cars, reducing gasoline use. Telecommuting, e-commerce, and online libraries and information sites significantly reduce the need for driving and flying and thus, the need for fuel. One fiber-optic cable, with about 150 pounds of silica, carries more messages than a ton of copper wire.⁶⁴

Digital storage of documents, data, photos, and so on, reduce the need for paper (and the

amount of trash produced.) Specific examples suggest the reductions: A large insurance company reduced its use of paper by 100 million pages in a nine-month period by storing its manuals digitally instead of printing them. A computerized system for recording insurance claims replaced more than 30 million index cards. We use email and texting instead of sending letters and cards on paper. Electronic payments eliminate paper bills and checks. We read books, newspapers, magazines, and so on, on tablets, e-readers, and smartphones, reducing paper use. The decline in business for the U.S. Postal Service and printed newspapers, while population and economic activity grow, are indications of these reductions. But do we actually use less paper than we did before? I could not find clear data for total paper use. However, between 2001 and 2011, annual consumption of newsprint for daily newspapers in the United States dropped by an estimated 61%, and the number of pieces of first class mail dropped by about 24%.⁶⁵

We take, post, and share far more photos (billions per month) than we did when we made prints and slides. This is one example of a phenomenon that occurs in many fields: as a product or service becomes more efficient and cheaper, we use more of it. It seems that people have certain levels of cost that they are willing to accept. We do more and use more when the cost goes down. Perhaps we have increased our use of resources as we have increased our use of computing technology. Certainly, we shift resources from areas where there are savings to other uses, including improved medical technology, more music and video, easier access to education, and other products and services that bring us benefits.

26. Records from 18th-century French villages showed that the median age of death was lower than the median age of marriage. Until recent generations, parents had to endure the deaths of most of their children. Starvation was common. In the United States, life expectancy at birth increased from 47.3 years in 1900 to 77.9 in 2007. Worldwide average life expectancy increased from approximately 30 in 1900 to approximately 64 in 2006. Science and technology (along with other factors such as education) reduced or almost eliminated typhoid, smallpox, dysentery, plagues, and malaria in most of the world. Deaths at work, during travel, and by accidents declined dramatically.⁶⁶

In the early 2000s, Americans spent less than 10% of family income on food, compared to approximately 47% in 1901. Agronomist Norman Borlaug, who won a Nobel Peace Prize for his work in improving agricultural productivity, reported that when new forms of wheat and crop management were introduced in India, yields rose from 12.3 million tons in 1965 to 73.5 million tons in 1999. In about the same timeframe, U.S. production of its 17 most important crops increased from 252 million tons to 596 million tons, but used 25 million fewer acres. Nicholas Eberstadt, an expert on population, reported that food supplies and gross domestic product have been growing faster than population for decades in most areas of the world, in both developing and developed countries.⁶⁷

The benefits of telecommunications and information technology are enormous in developing countries. A report of a United Nations Conference on Trade and Development, for example, observes that developing economies can make productivity gains worth billions of dollars by encouraging the growth of electronic commerce. The report said that “it is because the internet revolution is relevant not just to the high-tech, information-intensive sectors but also to the whole organisation of economic life that . . . developing countries stand a better chance of sharing in its benefits earlier than in previous technological revolutions.”⁶⁸

Technology is certainly not the only factor in improving quality of life. Progress against disease, discomfort, and early death depends on the stability, freedom, and flexibility of political and economic systems as well. Measuring quality of life is subjective, and some find other measures more important than the few we cited above. But, for many people, these data suggest that technology has contributed much to human well-being.

7.4 Making Decisions About Technology

No one voted for this technology or any of the various machines and processes that make it up.

—Kirkpatrick Sale⁶⁹

7.4.1 QUESTIONS

We saw, in Section 7.3, that the determination of what are true needs depends on our choice of values. Throughout this book, we saw controversies about specific products, services, and applications of computer technology (for example, personalized advertising, anonymous Web surfing, and face recognition systems). How should we make decisions about the basic question of whether to use a whole technology, or major segments of it, at all? Who would make such decisions?

Most people in science, engineering, and business accept, almost without question, the view that people can choose to use a technology for good or ill. Some critics of technology disagree. They argue that technologies are not “neutral.” Neil Postman says, “Once a technology is admitted [to our culture], it plays out its hand; it does what it is designed to do.”⁷⁰ This view sees the technologies themselves as being in control.

In the view of some critics of computing technology, big corporations and governments make decisions about uses of the technology without sufficient input or control by ordinary people. Kirkpatrick Sale’s lament at the beginning of this section expresses this view: there was never a vote on whether we should have computers and the Internet. Some people argue that we should not use a new technology at all until we have studied it, figured out its consequences, and made a determination that the consequences are acceptable. The idea is that if the technology does not meet certain criteria, we would not permit its development and use.

This view leads to a few basic questions. Can a society choose to have certain specific desirable modern inventions while prohibiting others or prohibiting whole technologies? How well can we predict the consequences of a new technology or application? Who would make the decisions? We consider the first question here and the others in the next few sections.

How finely can we make decisions about acceptable and unacceptable technologies? In response to a criticism that the tribal life he extolled would have no pianos, no violins, no telescope, no Mozart, Sale replied, “[I]f your clan thought that the violin was a useful and nonharmful tool, you could choose to invent that.”⁷¹ Perhaps critics of computing technology who recognize its value to disabled people would permit development of applications for them. The question is whether it is possible for a clan or society to choose to invent a violin or a book reader for blind people without the technological and economic base on which development of these products depends. That base includes the freedom to innovate, a large enough economy to get materials from distant sources, and a large number of potential applications that make the research, development, and production of the basic ingredients of these products economically feasible. It is unlikely that anyone would even think of developing a book reader for the blind if some of the components did not already exist in prior products (for example, perhaps, a photocopy machine).

Telemedicine: A bad application of technology?

In Chapter 1, we described long-distance medicine, or telemedicine, as a benefit of computer technology. Computer and communications networks make possible remote examination of patients and medical test results, and they make possible remotely controlled medical procedures. You should be able to think of potential privacy and safety problems with such systems. You might think of other objections as well. Should we ban telemedicine?

Several states passed laws prohibiting the practice of telemedicine by doctors who are not licensed in that state. The main argument they give for the laws is safety, or concern about out-of-state “quacks.” The laws will “keep out the charlatans and snake-oil salesmen,” according to one supporter.⁷² Also, telemedicine could increase the influence of large, well-financed

medical centers—to the detriment of local physicians in private practice. Large hospitals might become the “Wal-Mart of medicine,” says one writer. Telemedicine might make medical care even more impersonal than it is already.

Is concern for patients the real reason for the laws? The arguments about charlatans and quacks seem weak, considering that the laws target doctors who are licensed, but in another state. Many doctors who support the bans see telemedicine as a significant competitive threat. As the director of one state medical board put it, “They’re worried about protecting their turf.”⁷³ The laws restrict competition and protect established special interests—a risk of any mechanism designed to prohibit a new technology or product.

7.4.2 THE DIFFICULTY OF PREDICTION

A brief look at the development of communications and computer technology suggests the difficulty of evaluating the consequences and future applications of a new technology. Early computers were developed to calculate ballistics trajectories for the military. The PC was originally a tool for doing computation and writing documents. No one but a few visionaries imagined most of their current uses. Each new technology finds new and unexpected uses. When physicists began developing the World Wide Web, who would have predicted online auctions, social networking, or sharing home video? Would anyone have predicted even a small fraction of the ways we use smartphones? Postman’s statement that a technology does “what it is designed to do” ignores human responsibility and choice, innovation, discoveries of new uses, unexpected consequences, and social action to encourage or discourage specific applications. Computer scientist Peter Denning takes a different view: “Although a technology does not drive human beings to adopt new practices, it shapes the space of possibilities in which they can act: people are drawn to technologies that expand the space of their actions and relationships.”⁷⁴ Denning says people adopt technologies that give them more choices. Note that he does not say more choices of consumer products, but more actions and relationships. Don Norman also suggests that society influences the role of a technology when he says, “The failure to predict the computer revolution was the failure to understand how society would modify the original notion of a computational device into a useful tool for everyday activities.”⁷⁵

-
- *The telephone is so important, every city will need one!*
—Anonymous
 - *My personal desire would be to prohibit entirely the use of alternating currents. They are unnecessary as they are dangerous.*
—Thomas Edison, 1899
 - *I think there is a world market for maybe five computers.*
—Thomas J. Watson, chairman of IBM, 1943
 - *Computers in the future may . . . only weigh 1.5 tons.*
—*Popular Mechanics*, 1949
 - *There is no reason for any individual to have a computer in their home.*
—Ken Olson, president of Digital Equipment Corp., 1977
 - *The U.S. will have 220,000 computers by the year 2000.*
—Official forecast by RCA Corporation, 1966. The actual number was close to 100 million.
-

Figure 7.4 Predictions.⁷⁶

How well can a government committee, a think tank, or a computer industry executive predict the consequences of a new technology? The history of technology is full of wildly wrong predictions—some overly optimistic, some overly pessimistic. Consider the quotations in Figure 7.4. Some scientists were skeptical of air travel, space travel, and even railroads. (They believed that passengers would not be able to breathe on high-speed trains.) The quotations in Figure 7.4 reflect a lack of imagination about the myriad uses people would find for each new technology, about what the public would like, and about what they would pay for. They demonstrate humorously that many experts can be utterly wrong. We examine the prediction problem more seriously and in more depth by considering arguments made by computer scientist Joseph Weizenbaum in 1975 against the development of a particular computer technology: speech recognition systems.⁷⁷ We now have more than 35 years of hindsight. However, many inexpensive applications of speech recognition had already appeared by the early 1990s. Here are Weizenbaum's objections, accompanied by comments from our perspective today.

- *“The problem is so enormous that only the largest possible computers will ever be able to manage it.”* Speech recognition software runs on smartphones.
- *“. . . a speech-recognition machine is bound to be enormously expensive, . . . only governments and possibly a very few very large corporations will therefore be able to afford it.”* Millions of people own smartphones and other devices that include speech recognition.

- “*What can it possibly be used for?*” Much more than I will mention here. (Speech recognition technology is a multibillion-dollar industry.) We can search the Web from a cellphone by speaking what we want instead of typing. We can call a business, speak the name of the person we want to reach, and automatically connect to that person’s extension. Other customer-service applications include checking airline flight schedules, getting stock quotes and weather information, conducting banking transactions, and buying movie tickets on the phone by speaking naturally instead of pushing buttons.

Recall some of the applications described in Sections 1.2.3 and 1.2.4: training systems (e.g., for air traffic controllers and for foreign languages) and tools that help disabled people use computers and control appliances in their homes. People who suffer from repetitive strain injury use speech recognition input instead of a keyboard. IBM advertised speech-input software for poets, so they can concentrate on poetry instead of typing. People with dyslexia use speech recognition software so they can write by dictation.

Speech translation systems recognize speech and translate it into other languages. Full translation is still a difficult problem, but tourists, business people, social service workers, hotel reservations clerks, and many others use specialized versions.

Voice-activated, hands-free operation of cellphones, car stereos, and other appliances in automobiles eliminates some of the safety hazard of using these devices while driving.

- *The military planned to control weapons by voice command, “a long step toward a fully automated battlefield.”* Some argue that we should have the best possible weapons to defend ourselves. Others argue that, if wars are easier to fight, governments fight more of them. If countries fight wars with remotely controlled automated weapons and no humans on the battlefield, is that an improvement over wars in which people are slaughtered? What if only one side has the high-tech weapons? Would that cause more wars of aggression? Is there any technology that the military cannot or does not use? Should we decline to develop strong fabrics because the military can use them for uniforms? Clearly, military use of high-tech tools raises serious ethical and policy questions. Are these questions sufficient reason to abandon or condemn a technology?
- *Governments can use speech recognition to increase the efficiency and effectiveness of wiretapping.* Abuses of wiretapping concerned Weizenbaum (e.g., tapping done by oppressive governments). He does not explicitly mention wiretapping of criminal suspects. One can argue that governments can use the same tool beneficially in legal wiretapping of suspected criminals and terrorists, but it is true that speech recognition, like many other technological tools, can be a danger in the hands of governments. Protection from such abuses depends in part on the recognition

of the importance of strictly controlling government power and in part on the appropriate laws and enforcement mechanisms to do so.

Discussion of Weizenbaum's objections is important for several reasons. (1) Although Weizenbaum was an expert in artificial intelligence, of which speech recognition is a subfield, he was mistaken in his expectations about the costs and benefits. (2) His objections about military and government use highlight the dilemma: Should we decline to develop technologies that people can misuse, or should we develop the tools because of their beneficial uses, and use other means, including our votes and our voices, to influence government and military policy? (3) Weizenbaum's argument against development of a technology because of its expected cost is similar to arguments expressed by others about current and future computer applications and other technologies. For example, a common objection to some new medical technologies is that they are so expensive that only the rich will be able to afford them. This shortsighted view can result in the denial of benefits to the whole population. For many new inventions, prices are high at first but quickly come down. A computer chip developed to float on the retina of the eye and send visual signals to the brain has the potential to restore sight to some blind people. The initial cost was \$500,000. Should we ban it because it would be available only to the very rich? The developer of the chip expected the cost to come down to \$50 with mass production.

Weizenbaum was not trying to evaluate computer technology as a whole but was focusing on one specific application area. If we are to permit the government, or experts, or the people via a majority vote to prohibit development of certain technologies, it is essential at least that we be able to estimate the consequences—both risks and benefits—of the technology fairly accurately. We cannot do this. The experts cannot do it.

But what if a technology might threaten the survival of the human race? We consider such an example in the next section.

7.4.3 INTELLIGENT MACHINES AND SUPERINTELLIGENT HUMANS— OR THE END OF THE HUMAN RACE?

Prominent technologists such as Hans Moravec, Ray Kurzweil, and Vernor Vinge describe a not-very-distant future in which intelligence-enhancing devices, artificial intelligence, and intelligent robots change our society and our selves in profound ways.* The more optimistic scenarios include human use of intelligent machines and services of many kinds. People might acquire advanced mental powers through brain implants and computer–brain interfaces. When someone has a stroke, doctors might remove the damaged part of a brain and replace it with a chip that performs the lost functions, perhaps with a large amount of extra memory or a chip to access the Web directly. Why wait for a stroke? Once the technology is available, healthy people will likely buy and install such implants.

* I include some references at the end of the chapter.

MIT robotics researcher Rodney Brooks, for example, suggests that by 2020 we might have wireless Internet interfaces that doctors can implant in our heads. He says people might be just as comfortable with them as they are now getting laser eye surgery at a mall.⁷⁸ Will such implants make someone less human than a heart transplant or pacemaker does? What social problems will intelligence enhancement cause in the next few decades? What philosophical and ethical problems arise when we combine human and machine intelligence in such intimate ways?

Going farther into the future, will we “download” our brains to long-lasting robot bodies? If we do, will we still be human?

The technological singularity

The term *technological singularity* refers to the point at which artificial intelligence or some combined human–machine intelligence advances so far that we cannot comprehend what lies on the other side. It is plausible, says computer scientist Vinge, that “we can, in the fairly near future, create or become creatures who surpass humans in every intellectual and creative dimension. Events beyond such a singular event are as unimaginable to us as opera is to a flatworm.”⁷⁹

Some technologists welcome the idea of the human race transforming into an unrecognizable race of superintelligent, genetically engineered creatures within this century. Others find it horrifying—and others unlikely. Some see potential threats to the survival of the human race. They see the possibility of the machines themselves achieving human-level intelligence, then rapidly improving themselves to a superhuman level. Once robots can improve their design and build better robots, will they “outcompete” humans? Will they replace the human race, just as various species of animals displace others? And will it happen soon, say within the next 20 years or so?

Two estimates support these scenarios. One is an estimate of the computing power of the human brain. The other is based on Moore’s Law, the observation that the computing power of new microprocessors doubles roughly every 18 to 24 months. If the progress of hardware power continues at this rate, then by roughly 2030 computer hardware will be about as powerful as a human brain, sufficiently powerful to support the computation requirements of intelligent robots.

Both those who think an extreme advance in machine intelligence or human–machine intelligence is likely in the near future and those who criticize these ideas provide several reasons why it might not happen. Here are some of them. First, hardware progress might slow down. Second, we might not be able to develop the necessary software in the next few decades, or at all. Developments in AI, particularly in the area of general intelligence, have been much slower than researchers expected when the field began. Third, the estimates of the “hardware” computing power of the human brain (the sophistication of the computing power of neurons) might be drastically too low. Finally, some philosophers argue that robots programmed with AI software cannot duplicate the full capability of the human mind.

Responding to the threats of intelligent machines

Whether the singularity occurs within a few decades, or later, or not at all, many in the relevant fields foresee general-purpose intelligent machines within your lifetime. By its definition, we cannot prepare for the aftermath of the singularity, but we can prepare for more gradual developments. Many of the issues we explored in previous chapters are relevant to enhanced intelligence. Will software bugs or other malfunctions kill thousands of people? Will hackers hack brains? Will a large division open up between the superintelligent and the merely humanly intelligent? We saw that protections for safety and privacy in computer systems are often weak because they were not designed in from the start. It is valuable to think about potential problems of superintelligent systems and intelligence enhancement for humans well before they confront us so that we can design the best protections.

Bill Joy is cofounder of Sun Microsystems and a key developer of Berkeley Unix and the Java programming language. In his article “Why the Future Doesn’t Need Us,”⁸⁰ Joy describes his worries about robotics, genetic engineering, and nanotechnology. He observes that these technologies will be more dangerous than technologies of the 20th century (such as nuclear weapons) because they will be self-replicating and will not require rare and expensive raw materials and huge factories or laboratories. Joy foresees profound threats, including possibly the extinction of the human race.

What protections do people who fear for the future of the human race recommend? Joy describes and criticizes some before suggesting his own. Space enthusiasts suggest creating colonies in space. Joy observes that it will not happen soon enough. If it does, it might save the human race but not the vast majority of humans on earth. If colonists take the current technologies with them, the threat goes too. A second solution is to develop protections that can stop the dangerous technologies from getting out of control. Futurist Virginia Postrel suggests “a portfolio of resilient responses.”⁸¹ Joy argues that we could not develop “shields” in time, and if we could, they would necessarily be at least as dangerous as the technologies they are supposed to protect us against. Joy recommends “relinquishment,” by which he means we must “limit development of the technologies that are too dangerous, by limiting our pursuit of certain kinds of knowledge.” He cites, as earlier examples, treaties to limit development of certain kinds of weapons and the United States’s unilateral decision to abandon development of biological weapons. However, relinquishment has the same kinds of weaknesses Joy attributes to the approaches he rejects: they are “either undesirable or unachievable or both.” Enforcing relinquishment would be extraordinarily difficult, if not impossible. As Joy recognizes, intelligent robots and the other technologies that concern him have huge numbers of potentially beneficial applications, many of which will save lives and improve quality of life. At what point should governments stop pursuit of knowledge and development? Ethical professionals will refuse to participate in development of some AI applications, but they too face the difficult problem of where to draw the line. Suppose we develop the technology to a point where we get useful applications with legal and technological safety controls. How will

we prevent visionary or insane scientists, hackers, teenagers, aggressive governments, or terrorists from circumventing the controls and going beyond the prohibited level? Joy sees a relinquishment verification program on an unprecedented scale, in cyberspace and in physical facilities, with privacy, civil liberties, business autonomy, and free markets seriously curtailed. Thus, relinquishment means not only that we might lose development of innovative, beneficial products and services. We would lose many basic liberties as well.

Although we can find flaws with all proposals to protect against the dangers of powerful technologies, that does not mean we should ignore the risks. We need to choose appropriate elements from the various proposals and develop the best protections we can.

*Prediction is difficult, especially about the future.*⁸²

7.4.4 A FEW OBSERVATIONS

We have presented arguments against the view that we should evaluate and perhaps ban new technologies at the start. Does this mean that no one should make decisions about whether it is good to develop a particular application of a new technology? No. The arguments and examples suggest two things: (1) that we limit the scope of decisions about development of new technology, perhaps to particular products, and (2) that we decentralize the decision-making process and make it noncoercive, to reduce the impact of mistakes, avoid manipulation by entrenched companies who fear competition, and prevent violations of liberty. We cannot often predict the decisions and the results of decisions made by individual engineers, researchers, programmers, entrepreneurs, venture capitalists, customers, and teenagers who tinker in their garages, but they have a valuable robustness. The fundamental problem is not *what* decision to make about a particular technology. Rather, it is to select a decision-making *process* that is most likely to produce what people want, to work well despite the difficulty of predicting consequences, to respect the diversity of personal opinions about what constitutes a desirable life style, and to be relatively free of political manipulation.

When we consider the most extreme potential developments, such as superintelligent robots, what level of certainty of dire consequences should we require before restricting the freedom to develop technologies and products that might have marvelous benefits?



EXERCISES

Review Exercises

- 7.1 What is one significant criticism of Wikipedia?
- 7.2 What questions do we use to evaluate computer models?
- 7.3 What is one common use of Internet access in rural areas of developing countries?
- 7.4 Give one of the neo-Luddite criticisms of electronic commerce.
- 7.5 Give an example of a mistaken prediction made about computers.