THE NARRATIVE

JOSEPH MENN

"Hackers Live by Own Code"

About the Author: Joseph Menn is a journalist who specializes in cyber-crime issues. His book *Fatal System Error: The Hunt for the New Crime Lords Who Are Bringing Down the Internet* (2010) explores the links between hacking and international crime.

About the Article: This article considers some of the codes of ethics that hackers use to describe their own activities online and some of the various justifications they offer for their activities.

It wasn't Mary Ann Davidson's worst nightmare, but it was close.

A fax from a hacker in the Middle East landed on her desk at Oracle Corp., proclaiming the discovery of a hole in the company's database software through which he could steal crucial information from such customers as Boeing Co., Ford Motor Co. and the CIA. The fax warned Davidson, the company's chief security officer, to contact the hacker immediately—or else.

Luckily, the hacker hadn't found a real hole; he'd just misinterpreted a function of the program. More surprisingly, he meant no harm.

"The sort of threatening tone he took was really only to get our attention," Davidson said. "He actually turned out to be a nice guy."

The confrontational style of Davidson's hacker isn't unusual. As they troll through other people's computer networks, hackers abide by their own quirky rules of etiquette. What would strike most folks in corporate America as bad manners or worse may be considered the height of courtesy in hackerdom.

In large part, that disconnect stems from the fierce individualism of hackers—they are, after all, the sort of people who set aside the instruction manual and take a machine apart to see how it works. Though they inhabit a lawless domain where no data are considered private and "No Trespassing" signs are meaningless, they adhere to their own codes of ethics that vary depending largely on what motivates the hacker to hack.

Sometimes it's fame. Now and then it's money. Often it's a selfless desire to make software more secure. And occasionally it's a yearning to wreak senseless havoc.

The frequency of such attacks is on the rise, capped by the Blaster worm and SoBig virus that overpowered e-mail programs and crashed computer systems this summer. Computer Economics Inc. of Carlsbad, Calif., estimates that damage caused by hackers will cost companies and consumers $12.5 billion this year, up 13% from 2002.

Most hackers aren't malicious, security experts agree. But from afar, it can be difficult to distinguish the saboteurs from the merely curious, because they use the same tools, travel in the same virtual circles and often share a disdain for the rule of law.

Their philosophy predates personal computers, going back to the days when pranksters manipulated the telephone system to make free long-distance calls and cause other mischief. The personal rules that guide them today generally allow them to break laws, as long as they believe nobody will get hurt.

FIRMS ARE FAIR GAME

This maverick outlook is best personified by Kevin Mitnick, either the most notorious hacker or the most demonized, depending on your point of view. He stole millions of dollars' worth of software after cracking into

the computer systems of big companies such as Sun Microsystems Inc. and Motorola Inc. But he said he never sold any of it or otherwise profited from his electronic theft.

Mitnick, now 40, served five years in federal prison. Yet that hasn't deterred a younger generation of hackers who view private companies as fair game as long as no data are destroyed or profit turned. For many of them, hacking is just something their curiosity compels them to do.

Adrian Lamo, a 22-year-old hacker from Sacramento, always viewed his hacking habit as harmless at worst and helpful at best. If he has a chance to inform people about a security flaw in a company's internal network, he considers the disclosure a form of public service.

Lamo says he can't help it. He just starts wondering, then he looks for holes in a company's infrastructure, and he's in.

"When I'm curious about something, it's difficult to *not* seek out security problems," he said.

Working sporadically during long nights in Kinko's copy shops two years ago, Lamo used his battered Toshiba laptop computer to burrow deep into WorldCom Inc.'s internal networks. By the time he was done, he could have redirected the phone giant's employee paychecks to his own account or shut down the system of WorldCom customer Bank of America Corp.

Lamo did neither.

Instead, he recounted his exploits to a hacker turned journalist at SecurityFocus.com, a Web site devoted to tracking hacks, holes and fixes. SecurityFocus then called WorldCom executives and told them Lamo was happy to answer any of their questions. After Lamo showed WorldCom what he had done and how to prevent it from happening again, the company publicly thanked him for improving its security.

Part of Lamo's creed is a refusal to take financial advantage of anything he finds. The biggest compensation he's ever accepted from a company he's broken into, he said, was a bottle of water.

Chris Wysopal used to feel the same way when he worked at an outfit known as the Lopht, a band of security enthusiasts in a Boston apartment strewn with spare computer parts salvaged from area trash bins.

Claiming a dedication to telling software buyers the unvarnished truth, the Lopht crew published free security warnings on its Web site and in e-mail newsletters. Those warnings often were accompanied by programs to help people test whether their computers were vulnerable to attack.

In Wysopal's view, hacker etiquette didn't require him to give software makers advance warning before publishing his discoveries—even though his reports could aid the unscrupulous. Without the threat of public exposure and the fear that malicious hackers would use the newfound information, he figured, software makers wouldn't have incentive to make fixes in a timely manner.

"They dealt with security like a feature request—they would get around to it in the next version," Wysopal said.

The shaming tactics started working, so well that by 1999, Wysopal was forced to reconsider what constituted appropriate hacker behavior.

After the Lopht publicized a problem with a piece of Microsoft Corp. software for server computers, the company responded that it would have been happy to fix the mistake if only it had been given the chance. Instead, Microsoft had to race to develop a fix and get it to customers in time to head off an assault.

END TO FREE-FOR-ALL

Wysopal, along with a great number of his fellow hackers, realized the days of the free-for-all should end. It was no longer morally defensible to tell malicious teens how to hurt firms and their customers before they had the tools to defend themselves. Now he works with software makers to develop patches before blowing the whistle.

"It isn't as much fun," said Wysopal, who helped the Lopht morph into a computer security company called @stake Inc. "But if we publish right away, we are really arming the bad guys."

For other hackers, proper etiquette is dictated by the pursuit of money.

The most direct angle is simply to tell the software company there's a bug, then request a fee to explain it.

"If I come up with a vulnerability and I inform the source that I've discovered it, but I say, 'Would you mind paying me $5,000 to help you close it?' from my perspective that's a very reasonable request," said Bob Weiss, president of Password Crackers, Inc., in North Potomac, Md., which helps companies recover information hidden on their machines.

But what looks like a reasonable request to a hacker is often perceived as extortion by the company being asked to shell out. That's how one California software firm reacted after it heard from a hacker who had found a hole in its Web-messaging system and offered to explain it—for $10,000.

"The company got pretty mad," said Jennifer Granick, a cyber law specialist at Stanford University who represented the hacker in 2000. "It's very difficult for some cocky 18-year-old kid to approach a company without it feeling threatened." After Granick smoothed things over, the company agreed not to press charges.

There's also the loss-leader approach. After identifying a problem and explaining it, many hackers offer to look for additional glitches in exchange for a consulting fee.

Even that strategy backfired on a Boxboro, Mass., security group called SnoSoft. In 2002, SnoSoft researchers found a hole in a version of the Unix operating system made by Hewlett-Packard Co. The hackers told HP they would explain it for free, but they also asked to be paid for additional work.

"We made it clear we wouldn't charge [for the initial bug], because that would be extortion," SnoSoft co-founder Adriel Desautels said.

HP declined to offer SnoSoft a contract. Instead, the company threatened to sue under the Digital Millennium Copyright Act of 1998, which prohibits some attempts to tinker with programs to see how they work.

To computer security experts—including some inside HP—that threat amounted to a gross violation of etiquette on the part of HP. The company backed down and recently said it would never use the digital copyright law to stifle research. The Palo Alto computing giant declined to discuss the SnoSoft case.

For a few hackers, there is only one principle that matters: Do as much damage as possible.

That may have been the goal of a group of Chinese hackers who reverse-engineered a patch designed to fix a devastating hole in most versions of Microsoft's Windows operating system for PCs and servers. Within days, the hackers published a program to seize control of unsuspecting computers, which was used by others in the Blaster worm attack this summer.

COUNTERATTACKS INCREASE

With malevolent programs on the rise, large software companies are trying to get a handle on the problem. A consortium of software giants including Microsoft and Oracle has joined with security firms such as Symantec Corp. to formalize the etiquette of hacking so that software makers have time to patch holes before they are disclosed to the world at large.

The rules proposed by the new Organization for Internet Safety would give companies a month or so to develop and distribute a patch. Then another month is supposed to elapse before the hacker can disclose any details about the problem that the patch was designed to fix.

But hackers say they are unlikely to sign off on the rules, especially because they would neutralize the biggest weapon in their arsenal—the threat of public exposure.

In the meantime, companies that find themselves victimized by hackers are stepping up their counterattacks.

The New York Times wasn't amused when Lamo, the hacker who helped WorldCom beef up its network security, bragged to SecurityFocus that he had wriggled into the newspaper's computers.

Once inside, Lamo perused records of contributors to the paper's Op-Ed page (including the Social Security numbers and home phone numbers of former heads of state), conducted database searches using the paper's Lexis-Nexis account and added himself to a list of expert sources on hacking.

Unlike WorldCom, the New York Times called the FBI. In September, federal prosecutors in New York charged Lamo with the electronic equivalent of breaking and entering.

Out on bail, Lamo said he had no regrets about the way he hacked. "I always knew that the things I did could have consequences," he said.

Journal/Discussion Questions

1. What elements of the hacker's codes do you think are legitimate? Which ones do you reject? Why?

Living Together in Cyberspace
An Introduction to the Moral Issues

Introduction

If you are in your twenties or younger, you probably don't even remember a time when computers were not part of your life. In fact, for much of your life you have probably had ready access to the Internet. Today, the Web is everywhere, beginning with the phone that you are carrying. We are always connected, always on.

One of the consequences of growing up with ready access to the Internet is that we hardly realize the extent to which it has changed our lives. Nowhere is this more evident than in the realm of privacy. In previous

generations, individuals typically expected much of their lives to be private. This was, in effect, the default position. Now all individuals are much more likely to expect their lives to be continually on display.

MOORE'S LAW AND THE GROWTH OF COMPUTER TECHNOLOGY

Computer and information technologies have grown exponentially over the last two decades, and promise to continue to do so. Moore's law—developed by Gordon Moore, a co-founder of Intel—states (in everyday language) that computing power will double every 18 months, and this has proved to be a slightly conservative estimate. If you put $1,000 in the bank in the year 2000, and it grew at the same rate as computing power, you would have over $250,000 at the end of the year 2012, and you would break the $1,000,000 mark three years later. Moore's law is what allows you to buy a new computer every couple of years and get one that is twice as powerful as your old one, often for a little less money than you paid the last time. Of course, it is also what makes your current computer almost worthless on the used-computer market!

LIVING IN A COMPUTER-MEDIATED WORLD

The result of this tremendous growth in computing power has been the emergence of what I will call here a *computer-mediated world*. Relationships that had previously been direct, either directly to another person or to a physical object, have now become relationships that occur with the computer in between. Sometimes we may notice the presence of the computer, but increasingly the computer mediation slips into the background and remains unnoticed by most of us in everyday situations. Let me give you a few examples.

SHOPPING

Let's begin with shopping. Increasingly, many of us do our shopping online. I am certainly no exception to that rule. It is far easier to simply click on an item and have it delivered to my door a couple days later than to go to the store and go through the hassle of trying to find the right item at the right price. And this is, I think, a tremendous advantage: I'm able to save time and money and I am much more likely to get exactly the product that I want.

In the traditional mode of shopping, when I used to go to the stores themselves, my relationship to the objects that I might buy was either unmediated or mediated by a salesperson. Let me explain. If I'm walking around the store, picking up items, looking at them closely, perhaps trying them on if they are items of clothing, I am relating to those objects in a direct, unmediated way. If I walk into a higher priced store, it is more likely that I will be greeted by a salesperson who will show me possible objects or products to buy after I explain what I am looking for that day.

In online shopping, however, my relationship to the possible objects that I may purchase is mediated by a computer system. The computer presents possible objects for purchase and provides additional information about choices and models, colors, quantity, and the like. Yet it does so in a way that makes itself invisible. We simply see the pictures, the text, and the drop-down boxes and check boxes on the screen. Sometimes this computer mediation is a good thing; at other times, it may be of more benefit to the seller than to the consumer. In either case, it remains largely invisible.

Consider a single example of the situation in which this computer mediation may work to your disadvantage. Several years ago I was sending flowers to a cousin in another city for her birthday. There was a nice selection of flower arrangements from as low as $35 to as high as $200. Not wanting to appear to be totally cheap, I chose to send a nice arrangement that cost $60. When I returned to the same site a year later, I was surprised by the range of choices. Now floral arrangements stretched from $60 to $250. Had flowers gone up so much? No, the computer has simply recognized me from my previous purchase and displayed only those choices that it wanted me to consider. They still had floral arrangements at $35, but the cookie implanted on my machine told them that I was willing to pay $60, so they made sure that they did not show me the cheaper arrangements.

Of course, not all computer-mediated shopping works this way, and often by shopping online I am able to save money rather than spend more than I want. But it is important to note here that the computer mediation

largely disappears into the background, hiding itself from view. As savvy consumers, this is something that we should recognize and treat appropriately. One of the ethical issues here falls into the domain of business ethics: to what extent is this process of hiding certain items from view an acceptable business practice? This leads to another important ethical question: to what extent should such practices be regulated by the government in order to protect consumers from practices that are largely invisible to them?

Driving

Even something as simple as driving is now increasingly computer mediated. For years now, American cars have had antilock brakes. When you press on the brake pedal, the tiny computer devoted to this task translates your steady pressure into an intermittent pressure on the brakes that more effectively slows down the car without a risk of putting it into a skid. Recently in some new cars, cruise control has achieved a new dimension. Cruise control was already a computer-mediated way of driving for it automatically kept the car at a constant speed whether growing up- or downhill, accelerating or braking as needed. Now, in some cars it keeps track of the cars in front of you, slowing down your car when it gets too close to a vehicle ahead of you. Some cars will now alert you if you begin to drift out of your lane by voice or some type of vibration alert. A few cars even park themselves.

Navigation has become increasingly computer-dependent as well. Many cars come with built-in navigation systems, and for those without such systems, drivers often depend upon their cell phones for navigation directions. Again, what had previously been an unmediated relationship between driver and destination now becomes a mediated one, a relationship structured through a computer. In most cases this works very well. In cities that are not laid out on a grid, such as San Diego, the systems are particularly helpful. But the help comes at a price. As we increasingly depend upon the computer to give us directions, we become less skilled in orienting ourselves to the landscape. Moreover, anyone who has had the experience of following the commands of a navigation system when they are clearly in error knows how difficult it is to switch back to the mode of thinking for yourself. In a world of computer mediation, basic human skills often begin to atrophy.

Anyone who has seen Google's self-driving car—just search for the YouTube video about this—has had a glimpse of the future. This is a car that can drive itself in cities as well as the countryside for tens of thousands of miles without an accident. Imagine getting in your car, giving it a voice command for the destination, and sitting back and relaxing until you arrive at your chosen destination. Even though Americans may be hesitant to give up their sense of personal freedom associated with driving, it is easy to see the economic forces that will push us in this direction. First, fewer accidents will mean lower insurance rates, and this alone will establish a certain degree of pressure to move in this direction. Moreover, this is not simply a financial gain, it is also a *moral* gain. It is reasonable to imagine that as this technology becomes increasingly reliable, we may be morally obligated to use it because of the number of lives it will save. In 2009, there were almost 36,000 deaths related to motor vehicle accidents in the United States according to the U.S. Census Bureau. In other words, more people die every month in the United States by traffic accidents than died in the 9/11 attacks.

Other economic factors are going to contribute to this move toward computer-mediated driving. The reaction time of computer systems is far faster than human reaction time, and as a result it will be possible for computer-driven cars to follow much more closely than they would be able safely to do if the cars were driven by a human being. Imagine if we could simply double the number of cars in the given lane, the need to widen freeways would decrease dramatically, saving taxpayer dollars and preventing encroachment on private property by expanding freeways.

Social Relationships: Privacy, Intimacy, and Trust

Interpersonal relationships have become increasingly computer mediated as well. The advent and subsequent rise of Facebook, Twitter, and other social networking media has transformed the way human beings relate to one another. The result has been a world in which people are more connected than ever before, able to maintain friendships halfway around the globe without the barriers of time and space getting in the way. Moreover, with life online becoming the norm, we find that our notion of privacy has been transformed in ways that were unimaginable a generation ago. As Facebook and other social networking sites continue to fine-tune the

structures of privacy regulation in an online world, we find that there is simply less and less of an expectation of privacy in this new world.

This has left many people divided over the nature and function of privacy. Sometimes this is along generational lines, but not always. Defenders of traditional notions of privacy have advanced several interesting arguments, but the most intriguing of these relates to the connection between privacy, trust, and intimacy. The argument goes as follows. We have numerous relationships with people throughout our day that fall into the category of public relationships. However, most of us have a few people—at least one or two—with whom we have a special relationship. These are people we trust with information about ourselves that we would not necessarily divulge to the general public. This may be information about our deepest fears and hopes, about our vulnerabilities, and similar types of information about ourselves. None of this information is necessarily bad, but it is not intended for public consumption. Indeed, a key component of being a close friend to someone is precisely the willingness to share this type of information in a reciprocal fashion. Privacy establishes a zone of intimacy with a select number of people within which we can express those aspects of ourselves that are most vulnerable and sometimes most important. Indeed, these relationships of trust are typically our most important relationships. There is, in other words, a deep connection among privacy, intimacy, and trust, and if we value intimacy and trust in our lives, then privacy is necessary to protect such intimacy and trust.

Searching for Knowledge: The Ethics of Search Engines

One of the areas in which our lives have become increasingly mediated by computers is the search for knowledge. Search engines play an absolutely vital role in almost all aspects of our everyday lives now, but rarely do we stop and reflect upon these roles and upon the ethical constraints that should govern the operation of search engines. Search engines are, to put it metaphorically, the keys to the kingdom of knowledge, yet in a typical day we hardly notice the search engines we use. If they work well, and they usually do, they are as transparent as the clear glass in the window, revealing what lies beyond but concealing itself in the process. Moreover, the centrality of search engines grows day by day as our access to all kinds of knowledge shifts increasingly toward the World Wide Web and away from traditional print sources. Consider just a few of the ethical issues relating to this.

When you do a Web search on Google or some other search engine, you have undoubtedly learned that it doesn't pay to be coy: when searching for something, you need to be direct and specific. This has a very interesting consequence: if you look at the history of someone's Web searches during a day, you practically have a direct pipeline into that person's mind. Our search history tells us what we have been thinking about, and it also tells other people what we have been thinking about as well. Search histories are typically tracked closely by search engine companies, Internet providers, and anyone to whom they can sell this information, particularly advertisers. This information is gold to them, which is why they are happy to provide search capabilities to you for free. What they get in return is knowledge about your search habits, your preferences, your buying habits, and many other things of great interest to potential advertisers.

How private should this information be? Should there be government regulations limiting the uses to which it can be put? Is it sufficient simply to provide this information in the end-user license agreement that you sign but probably never read? What if your searches relate to potentially sensitive issues such as concerns about health, economic prospects, or even relationships? Continuing along the same lines, we can ask how long a search company should be allowed to retain such information about you. In some countries, they are mandated to destroy such information after a period of several months. In the United States, there are far fewer regulations governing the operations of search engines.

Or consider another ethical issue relating to search engines: which sites should come up first? This is not as easy a question to answer as you might at first think. The success of Google has been in its ability to match up your query with the information that you were searching for. The site that comes up first might not be the most popular, but in Google's very sophisticated opinion it is most likely to be the site that you are actually looking for. Indeed, Google makes no claim to the objectivity of its results, because there is nothing to be objective about. Their genius lies precisely in their ability to discern your subjectivity, that is, your subjective intent, and to match it up with the relevant Web sites.

You can see the possibility for abuse here, although I think there is absolutely no reason to suspect that Google in any way slants its search results. The possibility of abuse however lies in the fact that there simply is no objective standard by means of which the accuracy of the results can be measured. The search is successful simply if the search engine provides you with the information you want—or at least appears to do so. In the very early days of Internet searches, there was a disturbing blurring of the lines between paid advertisements and actual search results, but now search providers typically draw a clear distinction between the two types of results. However, we can see that these are murky waters, for there is no clear criterion for determining which site should come up first. Popularity is important, but the most popular sites are not necessarily the ones that you're looking for. The number of other sites that link to a particular site, what are called back links, is another important indicator, but that alone is hardly a guarantee that a particular site is the one that you are most interested in looking at. Typically, it is going to be some combination of factors—a complicated formula called an "algorithm"—that determines which sites come up first.

In fact, the problem was more complicated than this description suggests. These algorithms are highly sophisticated formulas for ranking Web sites, and as such they are closely guarded secrets that search companies are not going to divulge to the general public. In fact, there are two very good reasons why they should be kept secret. First, these are proprietary trade secrets that give one company an edge over its competitors, and they would be foolish to divulge such secrets to the general public. In doing so, they would lose their competitive edge. Second, it is important to keep these algorithms secret in order to prevent unscrupulous Web site operators from manipulating the formulas so that they can increase the rating of their sites. Major search firms continually guard against such manipulation, for this would undermine public trust in the reliability of the results.

Now we are in a position to recognize the complexity of this issue: search engines are crucial to the quest for knowledge, but by their very nature they cannot be transparent in their operation to the average Web user. As a result, we can never look "behind the screen," as it were, to see what's really going on and to make sure that the search results that we are viewing have not been manipulated. In other words, we just have to trust the Web search firms to present the results honestly.

Yet there is one more wrinkle in this entire fabric: increasingly, search engines have tended to tailor their results to particular users. For example, if I often read conservative columnists on the *New York Times* Web site, when I do a search on a particular topic in the political domain, search engines are now more likely to deliver conservative results to me before they display results from columnists at the other end of the political spectrum. Or, to take a less controversial issue, if my search engine knows that I am a Mac user—and it inevitably does know this—it may be more likely to yield results of a software search in such a way that priority is given to Mac software. All of this contributes presumably to making me more satisfied with my search results, but it does so at the cost of blurring even more fundamentally the criteria for what comes up first in a Web search.

These questions would be less important if it weren't for the fact that the Web is increasingly not simply our dominant source of information, but often our *only* source of information. At the very heart of our societal search for knowledge is a black box into which we cannot peer; it is the hidden engine that controls the whole process but whose rules are hidden from our view. Even for those of us who at a personal level have tremendous confidence in both the ability and integrity of search firms such as Google, this is a dangerous situation. Google, like many such firms, bears the imprint of its founders and has been shaped by their vision. Yet we have seen how time after time founding executives eventually retire. At first, their successors might be brought up from within the ranks of the company, but eventually they are replaced by outsiders who share neither understanding nor loyalty to the vision that the founders of the company originally had. At that point, the company becomes one more largely anonymous corporate entity, indistinguishable from hundreds of others—and no more trustworthy than those others.

The preceding remarks are simply a small sampler of the many, many ways in which we live in a computer-mediated world. This is a world that creates an array of new moral issues, issues that in some cases have never existed in the past, and issues to which in the future we will need to provide new answers. You will see some of these new issues as well as some of these new answers in the readings that follow in this chapter.

ETHICS IN A POLICY VACUUM

Given this exponential growth, it is not surprising that computing has transformed our world; nor is it surprising that it has brought about changes so quickly that it has outstripped the ethical rules that usually guide our decisions. Computer ethics arises in response to what James H. Moor, a noted computer ethicist at Dartmouth, has called a policy vacuum. Technology gets out ahead of us, and we have to work hard on our policies and ethical guidelines to catch up with the new technology. In the area of computing, some of the challenges are novel enough that, in attempting to address them, we may find ourselves transforming, not just developing, ethical theory.

Although the area of computer ethics covers a wide range of topics, many of the specific issues fall into one or more of several major categories: privacy, ownership of information, and security. In addition, there are important, more general ethical questions about the ways in which computers and the Internet have transformed our lives. Some of these issues are treated in James Moor's article, reprinted here, "Should We Let Computers Get under Our Skins?" Before looking at those, let's consider some of the specific moral problems raised by these technological advances.

PRIVACY AND CONTROL OF PERSONAL INFORMATION

The distinction between public and private is central to the American democratic tradition, for the private domain has long been considered the domain within which individual freedom is most fully exercised. In particular, the private has long been considered comparatively safe from government scrutiny.

BIGBROTHER.GOV OR BIGBROTHER.COM?

Computers have changed all that, and they have done so in a very interesting way. Various computer-based technologies allow governments, corporations, and even individuals to collect an amazing amount of information about people. Imagine all the ways in which an average person leaves electronic footprints during the day. Using computers at home and work, individuals leave a vast amount of information about themselves, their reading and buying habits, their business dealings, their personal likes and dislikes, and so forth. Prior to the advent of computers, one person could perhaps have followed another person around, noting the other person's actions, contacts, and the like. The rise of computers now allows a single individual to track thousands of people simultaneously, finding a range of details about them that would previously been impossible.

In the UK, the use of closed-circuit television (CCTV) cameras is more extensive than anywhere else in the world. Part of this is due to the fact that in the UK there are fewer constitutional protections of privacy, and another part is due to the fact that they endured decades of terrorist attacks on their own soil by the IRA. The result has been an extensive network of CCTV cameras not only in the major cities, but even in the small villages in the UK. Typically, the average resident appears on dozens, if not hundreds, of CCTV cameras on any given day. Moreover, video cameras that read and record the license plate numbers are now in place on all British roads, so it is a simple matter to keep track of a given car's movements throughout a specific day. This type of surveillance was already well established by 2010, and this massive use of video surveillance promises to become even more pervasive and finely tuned for the 2012 Olympic Games in London. In some cities, British authorities have already begun using two-way surveillance that allows them over a loudspeaker to speak to a potential lawbreaker. For example, if you toss an empty beer bottle on the ground in the UK, don't be surprised if you hear a voice asking you to please pick it up and dispose of it in the proper trash bin. Being under surveillance has simply become a part of everyday life in the UK.

What limits, if any, ought to be placed on the power of individuals, corporations, and governments to engage in such surveillance? Often we see doomsday scenarios in which government—a virtual Big Brother—controls the lives of ordinary citizens through increasingly extensive data collection. Since the 9/11 attacks, there have been growing political pressures in the United States to combat terrorism, often at the price of removing restrictions to the government's acquisition of information. However, this is not an issue limited solely to governments. Private corporations

often have access to information that could have a profound effect on the lives of individuals. Insurance companies and HMOs have been ethical hotspots for the privacy and control of information in recent years, and this has been a particularly thorny area insofar as insurance and health care are tied to employment. To what extent, if any, can insurance companies and employers and others share information about individuals with whom they deal? As the ability to share information increases, the need for clear and enforceable guidelines in this area increases as well.

Consider the following example. In Lancaster, Pennsylvania, local citizens got together and privately established a CCTV surveillance system that covers the entire town. They established a central video monitoring area, hired staff, and now the entire town is under twenty-four hour a day real-time video surveillance and the video feeds are then archived. Interestingly, there is nothing illegal about this. None of this was done by the government or with government funds. Those who monitor the cameras are not shorn peace officers, but simply employees of a private company.

In some cities, private companies have opened small, boutique-like shops that sell a variety of items. As is the case with most stores, there are small signs that indicate that there is video surveillance occurring, but this hardly tells the whole story. The purpose of the stores is not to sell products, although they do this, but rather to closely observe the customers to see what things they find interesting and what things they don't. Thus customers' online movements and other behaviors are closely tracked for research purposes in order to determine which packaging they find most interesting, which displays draw their attention, and the like.

FREE SPEECH, PRIVACY, AND CENSORSHIP

Consider some of the thorny issues relating to pornography. In the United States, pornography is defined in terms of the violation of local standards. But what counts as local? Is it the standards of the community of the person viewing the pornography? Or the standards of the community in which the purveyor of the pornography resides? What responsibility does the Internet Service Provider (ISP) have in hosting such a site? What restrictions, if any, ought to be placed on pornography sites?

What rights do individuals have to keep information about themselves off the Internet? In one case, critics of a police chief started putting personal information about him on the Web—pictures of his house, his address and phone number, and anything else they could find. Does the police chief have a right to privacy? Or, to take another example, in some cases anti-abortion activists have posted similar information about abortion providers. Again, to what extent—if any—do individuals have a right to privacy?

Closely related to this issue is the problem of data mining. The tremendous growth of computer power and storage, coupled with the vast increase in the amount of information stored in databases, has opened the door to powerful database searches that reveal more about individuals than they would have believed possible. Think of the many ways in which average individuals leave information about themselves in computers during a typical day: all swipes of cards for credit or identification purposes, preferred customer cards at the grocery store, all use of any computer they are logged onto, all telephone calls made or received—not to mention the countless surveillance cameras that may have picked up their images during the day. If all these pieces can be put together (and this is what advances in computing power and storage make possible), then an amazingly detailed picture of an individual's day begins to emerge.

Encryption plays an important role in this discussion as well. To what extent should individuals be allowed to encrypt their communications? Some levels of encryption (e.g., 128 bit) are virtually impossible to break. If criminals and terrorists can use such encryption, they can effectively prevent the government from monitoring their communications. What restrictions, if any, should the government be allowed to impose on encryption programs?

PROPERTY RIGHTS AND INTELLECTUAL PROPERTY

The rapid rise of file sharing Websites was the most dramatic signal of a change ushered in by the computer revolution: because digital media can be shared so easily and cheaply and without diminishing the original and without loss of quality, digitized property—most notably, music and video—began to speed around the Internet with an undreamed-of volume and speed.

As this happened, the traditional notion of private property—based on the paradigm of land and other physical objects—and property rights came under serious challenge. Our understanding of property rights was based on this traditional model, one in which giving a piece of property to someone else meant that you yourself lost something. For example, if I give a book to my friend, I no longer possess that book. I could, of course, xerox a copy of the book, but that would be a lot of work and the copy would not be as good as the original. In the digital world, however, this is no longer the case. If I have an electronic copy of a book (or a piece of music), I can send that to *all* of my friends and still not lose anything, and the copy they receive will be as good as mine.

Intellectual property is not exactly like a physical object, but we still have laws and policies based on the old paradigm of physical objects and there is significant disagreement about where to go from here. On the one hand, many maintain that the traditional laws of property ought to continue to apply unchanged and see all attempts to alter the law as a threat. On the other hand, some in the Internet community are fervently committed to the idea that all information ought to be freely available—and that means, among other things, removing restrictions on copying and duplicating materials, including music and video.

RESPONSIBILITY

In the age of the Internet, responsibility has become more diffuse in several ways. First, the Web allows the possibility of anonymity (or at least the illusion of anonymity), and this makes it easier for people to deny responsibility for their actions, even to themselves. Online pornography and Internet plagiarism are but two of the most obvious examples of the way in which the anonymity of the Web allows people to engage in behavior that, at least often, they would not perform in the public world of everyday life.

Moreover, the lines of responsibility in such cases are complex. What, for example, is the responsibility of ISPs for the actions of their clients? Do they, for example, have a responsibility to protect customers from cyber stalkers? From unwanted pornography? From unwanted advertisements? Or, to take another example, what responsibilities do college network administrators have if, for example, students are using college Internet connections to download pirated music and video? To order plagiarized term papers?

Responsibility issues have been transformed in another way as well. Increasingly, computer systems perform actions, and they do so with such speed that no human being can adequately monitor them. Let's look at an example from the stock market.

COMPUTERS AND THE STOCK MARKET

One of the areas in which we see computers as permeating the work process is the stock market. In 1998, the Securities and Exchange Commission opened the market to electronic exchanges, ostensibly to allow the average citizen with a desktop computer to compete. Soon, however, small traders dwindled into insignificance (and often poverty) when major traders began computerized trading on a massive scale. Let's take a closer look.

The Flash Crash of 2010

On May 6, 2010, the day at the New York Stock Exchange looked routine until about 3:00 in the afternoon. Stocks were down about 180 points, but that was not unexpected, given fears about the Greek debt crisis at that time. Suddenly, over the next twenty minutes, the Dow dropped over 1000 points in a period of a few minutes—slicing about one trillion dollars off the total value of the U.S. stock market—and then soared back up, recovering much of what had been lost. The precise cause of the crash remains unclear, a disturbing enough fact, but it seems to have been an unanticipated interaction among different exchanges with rules to slow down trading kicking in at different moments.

This was possible because much of the trading on the stock market is now done with computers. Indeed, much of it never even goes through the New York Stock Exchange (NYSE) any longer, thus avoiding many regulatory restrictions that might apply to the NYSE.

High-frequency Trading

At the beginning of the twenty-first century, stock trading underwent a radical transformation. Literally for centuries, stock trading was a human endeavor with actual human beings on the trading floor, buying and selling sometimes huge quantities of stocks back and forth. A giant billboard—the "Big Board"—kept track of current prices for buying and selling particular stocks. The most prestigious of these stock exchanges was the New York Stock Exchange, and only a limited number of stocks were traded on that exchange and only a limited number of traders were allowed on the floor. Essentially, it was unchanged except for minor upgrades to the hardware of the Big Board for many decades.

The advent of computerized trading changed all that. The brightest young mathematicians no longer went into research in science or into teaching at universities; they went to Wall Street, where they became known as "the Quants." (There are numerous books both by and about the quants.) They developed computerized trading formulas or algorithms ("algos" for short) that allowed them to discover tiny pockets of profit within a mountain of data and to quickly exploit the opportunities that thus emerge. One regulator has dubbed those who engage in such high-frequency trading "the Cheetahs," for they are both fast and fearless.

High-frequency trading now accounts for over sixty percent of the stock transactions every day in the United States. This is a realm in which humans cannot compete with computers except that humans write the algorithms upon which they are based. These algorithms allow computerized trading systems to purchase thousands of shares of stock in a nanosecond and sell them again a fraction of a second later, reaping a very small profit per-share that mushrooms into a much more substantial profit when we add in the number of shares traded and the number of trades that occur per minute. This is a world in which nanoseconds matter, that is, a world in which trading transactions are measured in billions of a second. Before you can blink, the computer may have made dozens of transactions, buying and selling again almost instantaneously, picking up small pockets of profit that would have been indiscernible to the human eye and acting on those opportunities with the speed that far outpaces human deliberation. This type of trading began with supercomputers early in the twenty-first century, and by 2009 was being noticed by Securities and Exchange Commission (SEC) regulators, for such lightning fast massive trades threatened to upset the market by what are called "flash crashes." Moreover, regulators began to get nervous when high-frequency traders began to move away from the established exchanges and to trade instead on what are called "dark pools," unregulated electronic trading exchanges in which high-frequency computerized trading is the norm.

One of the ethical concerns that regulators have expressed is that the rise of computerized high-frequency trading has left all others in the dust. This not only includes the average amateur day trader at home on a PC, but also includes many of the professional traders in established financial houses around the country. High-frequency trading depends on massive supercomputers, often located in buildings quite close to the exchange computers, to shave a few nanoseconds off the transmission time needed to execute trades, enabling such quick buying and selling that no one else can even keep up and track their progress. Moreover, loopholes in market regulation allow such supercomputers slightly faster access to market data than is given to other traders, with the result that these supercomputer trading systems are already a fraction of a second ahead of their competitors. The result, in the eyes of many, is a playing field that is no longer level, one that gives an unfair advantage to the high frequency supercomputers.

The 1987 Flash Crash

This was not, by the way, the only time such computer-driven trading had driven down the Dow. Back on October 19, 1987, computer-drive trading programs pushed the Dow to its knees, causing a 22.6 percent loss of value. It was the biggest percentage loss suffered on a single day by the Dow. Again, due to computers. That time it was "portfolio insurance" trading programs. Under pressure from regulators, new safeguards were introduced to slow the cascading effect of automated trading, but these were keyed to the time of day. There was less protection at 2:00 PM, and after 2:30 PM there was no protection against the cascade effect. The May 6, 2010, tumble occurred at 2:45 PM.

This is just one of many ways in which ethical issues are deeply interwoven into the ways in which computers are used in our world today.

The Diffusion of Responsibility

Computer systems often make decisions for us in daily life. Some of these decisions are relatively minor, but in some cases they highly significant. In some instances, computers may even decide when a missile is to be fired, especially in situations when an immediate response is crucial. In naval warfare, for example, where opposing ships are only a few miles apart and response time is correspondingly brief, we can imagine scenarios in which a ship's computerized fire system might retaliate on its own. In Iraq and Afghanistan, we have had little robots—dubbed "R2D2" by grateful soldiers—that fire an incoming mortar rounds. Since they need to respond immediately to incoming mortar attacks, they do not require a human operator to authorize each response. To whom do we attribute responsibility when something goes wrong? In the past, designers have built in a moment at which a human being must confirm the decision, but this is increasingly impractical. Even when a human being is involved, that person is often depending on information supplied by the computer. We can imagine a human being "in the loop" and necessary to confirm a decision to fire a missile, but the human being may be receiving information from computer-mediated sources. As computer-driven systems become increasingly autonomous, responsibility becomes more and more elusive.

The Digital Divide

This is not the only pervasive issue. While many (especially in education) find their lives permeated by computers in many ways, others are effectively barred from access to computers, primarily for economic reasons. The world is splitting into the haves and the have-nots, based in significant measure on access to computing power. Young children from middle and upper class families grow up comfortable with computers, happy to use them and confident of their own ability to figure out and solve problems. In an increasingly computerized world, they have an immediate advantage over children who do not grow up with such skills. Similarly, significant portions of the industrialized world are highly computerized, whereas many parts of the developing world lack even the basic necessities of life, much less access to computers.

The Arguments

Luciano Floridi

"*The Ethical Evaluation of WikiLeaks*"

About the author: Luciano Floridi is a contemporary philosopher whose work has defined the field of the philosophy of information. After a number of years teaching at Oxford, Floridi now holds a Research Chair at the University of Hertfordshire and the UNESCO Chair in Information and Computer Ethics. His most recent book, *The Fourth Revolution: The Impact of Information and Communication Technologies on Our Lives* (New York: Oxford University Press, forthcoming) deals with the broad impact of the information revolution. His *Philosophy of Information* was published by Oxford in 2011.

About the article: The publication of classified government documents by WikiLeaks under the leadership of Julian Assange raised a number of ethical issues, both about the dubious practices they sometimes revealed and about Assange's practice of publishing such documents. In this piece, Luciana Floridi examines some of the ethical issues underlying Assange's project.

As You Read, Consider This:

1. What are the two questions that Floridi focuses on in beginning of his analysis?
2. Explain what Floridi means by saying that confidential communication is a three-player game.
3. Whistleblowing, Floridi suggests, establishes a new, meta-game. Explain what he means by this.

Luciano Floridi, "The Ethics of WikiLeaks," Reprinted with the kind permission of the author. This is available online at http://philosophyofinformation.net/Blog/Entries/2011/1/5_Travels_through_the_east.html.

READING

4. How does the new meta-game destroy the relationship of confidentiality?
5. What happens to the phenomenon of accountability in the new meta-game?
6. What does Floridi mean when he says that information liberation arguments are not universalisable?

The Wikileaks phenomenon is intricate, but suppose we reduce its ethical evaluation to two questions: is whistleblowing ethical, even when motivated by resentment and the desire to harm its target? And is the facilitation of whistleblowing ethical, even if it might put at risk innocent people? A deontologist, convinced that telling the truth and never lying is an absolute must, is likely to appreciate whistleblowing as the right thing to do, independently of the reasons behind it. And a consequentialist may support Wikileaks as a means to maximise the welfare of the largest number of people, especially if risks are minimized by censuring sensitive information. So current answers in the mass media seem to converge: Wikileaks is a good thing. I am not entirely convinced.

Confidential communication is a three-player game—sender, receiver and referent—in which sender and receiver trust each other. The receiver, not the referent, trusts and holds responsible the sender for the truth of what is communicated about the referent. The referent may know about such communication and may even easily guess its contents (imagine a letter of reference), but there is confidentiality only if the receiver, not the referent, has access to the information exchanged. Accountability is present and connects sender and receiver.

Whistleblowing disrupts such a three-player game. In the new, metagame the sender is the whistleblower through Wikileaks, the whole world is the potential receiver, and the referents are the players in the previously confidential communication. This is problematic. The relation of confidentiality of the original game is destroyed: the new referents are also among the new world-receiver, which now holds the old sender responsible for what is communicated, not only for its truth (if you say that the moon is made of blue cheese, that is false, if I report that you said so then what I say is true). The metagame reinstates, somewhat hypocritically, the same rules it criticises: Wikileaks, quite rightly but inconsistently, defends the anonymity and confidentiality of its sources, which are likely to make an exception about the information transparency of their own identity, frowning upon MetaWikileaks, with leaks on leaks. Finally, the relation of accountability is missing. In the metagame, the whistleblower and Wikileaks might be good-willed and well-intentioned but are not bounded by professional codes of conduct or legal requirements. So the receiver, which is also the referent, is at the mercy of the sender. Wikileaks knows this and that is why it "whitemails" the world, i.e., it blackmails it by threatening to disclose even more damaging information through its "insurance file", should anything happens to Wikileaks or its spokesman Julian Assange.

Wikileaks itself shows that, without confidential communication, there would often be no communication at all. Thus, any argument in favour of Wikileaks to the effect that most of the information was already public or suspected anyway misses the point, which is that Wikileaks may undermine the possibility of future frank communication. Imagine an Academic Wikileaks that regularly publishes confidential information about the assessment of grants, the evaluation of book proposals, the reviewing of journal submissions, letters of reference for candidates and so forth. After the initial embarrassment, the whole system would come to a standstill.

Finally, "information liberation" arguments are not universalisable. The new Wikileaks' About file[1] holds that "publishing improves transparency, and this transparency creates a better society for all people". Yet this is naïve at best. First, because the value of information is not absolute, but relative to its use. Judas' kiss tells the truth about the identity of the kissed, but it hardly creates a better society. And second, because the value of the use of information is not absolute either, but relative to the goals that one is seeking to achieve, and the sort of possible world that one is trying to bring about. This is why personal details about religious and sexual orientations must be protected. Information "macht frei", but also doubles as a necessary condition for discrimination.

The lesson is simple: facilitating whistleblowing is morally good not absolutely, but only if the whistleblowing itself is morally good; and the latter is morally good not absolutely, but only if the specific cause it fosters is morally good. So the two conditionals call for an explicit, ethical commitment. And Wikileaks old About file[2] acknowledged this much: "Our primary interest is in exposing oppressive regimes in Asia, the former Soviet bloc, Sub-Saharan Africa and the Middle East, but we also expect to be of assistance to people of all regions who wish to reveal unethical behavior in their governments and corporations. We aim for maximum political impact." Unfortunately, this strong and explicit ethical statement has disappeared (Wikileaks is not a Wiki, so

old versions are no longer available from its website). Luckily, so far Wikileaks has picked up causes judged by most morally good. Support for Wikileaks would quickly vanish if the leaks undermined a cause such as the democratic movement in China. Yet the real ethical debate must concern the moral value of the causes supported by Wikileaks. And the concern remains: those who defend accountability should themselves be accountable. Who will whistleblow the whistleblowers, if their behaviour will become unethical?

[1]retrieved 12.12.10, http://www.wikileaks.ch/about.html
[2]archived 10.03.08, http://web.archive.org/web/20080314204422/www.wikileaks.org/wiki/Wikileaks:About

Journal/Discussion Questions

1. Floridi suggests that whistleblowers cannot universalize their position. Explain what is meant by this claim. What philosopher is Floridi implicitly referring to here?
2. Discuss the difference between deontological and consequentialist analyses of whistleblowing in general and WikiLeaks in particular.
3. Floridi ends his essay with a question: who will blow the whistle on the whistleblowers? How would you answer this question? What consequences flow from your answer to this question, especially in regard to privacy and censorship?

READING

JAMES H. MOOR

"*Should We Let Computers Get under Our Skins?*"

About the **Author:** James Moor is a Professor of philosophy at Dartmouth College, he is a primary figure in the growing area of computer ethics. His award winning article, "What is Computer Ethics?" is widely reprinted and regarded as a milestone for the study of computer ethics. He was an early pioneer in computer-assisted instruction in logic, including work on Carnap, Venn, and Proof Designer. He is co-author of *The Logic Book* and has written widely on the philosophy of artificial intelligence. Most recently, he is co-editor of *The Digital Phoenix: How Computers Are Changing Philosophy*.

About the **Article:** In this article, Moor looks at the question of whether we are gradually becoming cyborgs—part human, part computer—and what ethical limits ought to be imposed on this.

As You Read, Consider This:

1. What are cyborgs? In what ways are we already moving toward becoming cyborgs?
2. Explain the therapy/enhancement distinction. What is its significance in this article? What criticisms of this distinction does Moor consider?
3. Explain what Moor means by "the Borg argument."
4. What, according to Moor, are the three main areas in which we should be particularly sensitive to the coming of age of cyborgs?

Being connected with the passions also, the moral virtues must belong to our composite nature; and the virtues of our composite nature are human; so, therefore, are the life and the happiness which correspond to these.

—*Aristotle*

Being a human was ok, I even enjoyed some of it. But being a Cyborg has a lot more to offer.

—*Kevin Warwick*

READING

THE CASE FOR BECOMING A CYBORG

Aristotle suggests that human nature is fixed. Our human intellectual and moral virtues depend on our having this nature. If we changed our nature, we would change our virtues (excellences). Aristotle believed that if a friend became a god, for example, friendship with that person would cease because a god has a different nature than a human being. In the wake of evolutionary theory and modern genetics, the claim that human nature is fixed is not very plausible, but Aristotle's belief that shifting human nature might well alter moral virtue remains defensible. In today's scientifically changing world we need to confront this issue: if we change the kind of thing we are, what will be the consequences for ethics?

We will change ourselves genetically and we will change ourselves computationally as well. We will become cybernetic organisms—cyborgs—part human, part computer. The logical malleability of computers will allow us to go beyond what can be accomplished through genetic manipulation alone. The human body is the ultimate platform from which to launch new computer applications. It is likely that in the coming decades more and more computer hardware and software will be embedded in us. To what extent it should happen is the ultimate question, of course, but certainly there will be increasing pressure to produce cyborgs. Today the rationale and technology already coexist. First, we humans as creative creatures continually seek new ways to perform routine and not so routine tasks. Not infrequently our creative task solving involves the development of new tools. Second, the computer is the best master tool we have. The general-purpose computer is a meta-tool, a tool for making tools. If we have a task to do and we can express the task in terms of an appropriate algorithm to connect inputs to appropriate outputs, then in principle a computer can do it. In fact, even if we do not know an appropriate algorithm, computers using neural nets or genetic algorithms can sometimes evolve satisfactory computational structures for us. Third, considerable knowledge has been gained in recent years about interfaces between computers and living systems. We know that organic and inorganic structures can effectively interact at many levels—the organism level, the organ level, and the cell level. Someday nanomachines may interact in our bodies at the atomic level. Given that naturally curious humans love to find better solutions for problems, have a great master tool (the computer), and possess the perfect location (the human body) on which to store and operate the new devices, the gradual transformation of many humans into cyborgs, humans with computer parts, is all but certain.

Simply forbidding the implantation of computer chips because they are not natural, only artifacts, is not a plausible policy. This overly broad approach would not only prevent the use of beneficial computer implants but would rule out beneficial noncomputer implants such as artificial hip joints and dental crowns. Still, the thought of becoming a cyborg may seem rather repulsive. Who would want to have computer parts implanted? To become part computer? The idea of having a computer implanted may seem unnatural, possibly even grotesque, or at least something that undermines human dignity. But such a negative reaction is not defensible upon close examination. In fact, the transformation of humans into cyborgs has been taking place with no loss of dignity for years, although we do not commonly think of it in those terms.

For example, hundreds of thousands of people have cardiac pacemakers and defibrillators implanted to maintain regular heartbeats and heart rhythms (Lu, Anderson, and Steinhaus, 1995, Pinski and Trohman, 2000). Such implants not only promote life but also the quality of life. Totally implantable pacemakers have been in use since 1960, and programmable pacemakers were developed in the mid-1970s. The newest pacemakers can communicate via phone and the Internet. A patient needs only to wave a wand over his chest to pick up signals from the generator and then plug the wand into the phone line to send his physician an update on how the device and patient are doing. It would be hard to raise a principled objection against such beneficial devices. Implanting computerized cardiac devices is no more unnatural than putting other products of technology, like medicines or processed food, into our bodies. Given the alternatives for the cardiac patient, these vital, portable computer implants considerably enhance human dignity, not reduce it (Ocampo, 2000).

A similar case can be made for the benefits of implanting computer chips for vision. There are various projects underway to develop bionic eyes that will restore some level of vision to blind patients. Some approaches put chips on or under the retina and others connect computer chips to other parts of the visual system. In 2002, a Canadian farmer, who had lost his sight eighteen years earlier, had a bionic implant. A digital camera mounted on his glasses sent an image to a computer worn on his belt. The image was processed and sent to electrodes implanted in his

visual cortex. His vision was not fully restored, but he was not totally blind anymore. He was able to see well enough to navigate through rooms and even drive a car to a limited extent (Gupta, and Petersen, 2002).

Some diseases, such as retinitis pigmentosa or age-related macular degeneration, damage the rods and cones in retinas but leave the rest of their visual wiring, the ganglia cells that process information from the rods and cones and the optic nerve, intact. The various visual bionics under development hold great promise for bypassing the damaged areas of the visual system and restoring vision to the patient. In the United States, over a million people are legally blind and worldwide millions more. These cutting-edge bionic implants will offer enormous benefit.

More examples of beneficial computer implants can be marshaled, but I believe the case for the benefits of some computer implants is established. The debate is not whether humans should ever become cyborgs because in some cases, particularly where beneficial implants help overcome severe disabilities, justification for becoming a cyborg is clear. The transformation of some humans into cyborgs will continue to happen and it should. People who wish to have such helpful computer implants should be allowed to have them.

But how far should the conversion of humans into cyborgs go? What are the ethical boundaries? Could we find ourselves in a position that we would want to get away from computer implants but couldn't? Could the implants be used to track us? Reduce our autonomy? Give us freakish powers? In this paper I want to explore some potential ethical pitfalls of computer implants as well as the impact that computer implants might have on ethical theory itself.

Therapy versus Enhancement

The therapy versus enhancement distinction suggests a basis for a policy that would limit unnecessary computer implants. Given that the human body has natural functions, it might be argued that implanting chips in the body is acceptable as long as such implants maintain or restore the body's natural functions. In this spirit, consider the remarks of Michael Dertrouzos, a director of the MIT Laboratory for Computer Science, regarding the possibility of implants connected to the brain:

> Even if it would someday be possible to convey such higher-level information to the brain—and that is a huge technical "if"—we should not do it. Bringing light impulses to the visual cortex of a blind person would justify such an intrusion, but unnecessarily tapping into the brain is a violation of our bodies, of nature, and for many, of God's design. (Dertrouzos, 1997, p. 77)

The distinction between therapeutic applications and enhancing applications offers a criterion for limiting computer implants. Under this policy, pacemakers, defibrillators, and bionic eyes that maintain and restore natural bodily functions are acceptable. But giving patients additional pairs of robotic arms or infrared vision would be prohibited. It would endorse the use of a chip that reduced dyslexia but would forbid the implanting of a deep blue chip for superior chess play. It would permit a chip implant to assist the memory of Alzheimer's patients but would not license the implanting of a miniature digital camera that would record and play back what the implantee had just seen and heard. In a later book Dertrouzos stresses his point again, "Few people would implant a chip into their brain for less than life-and-death reasons. We have wisely set a high threshold for tampering with the core of our being, not just because of fear, but because of natural, moral, and spiritual beliefs" (Dertrouzos, 2001, p. 46).

Of course, even therapeutic applications raise ethical questions if they are not safe and effective or the patient has not given informed consent. But, let us assume safety, effectiveness, and informed consent. Does the therapy/enhancement distinction give a proper limit to computer implants? Although this policy generally forbids unethical implants, I believe it is too conservative and cannot be defended.

First Objection: Unclear Distinction

The line between therapy and enhancement is not always agreed upon. Consider the example of cochlear implants (Spelman, 1999). A microphone is worn behind the ear and a microcomputer filters and analyzes the signal from the microphone converting the sound into digital signals. These signals are sent by radio waves to a

READING

receiver implanted under the skin, then via a wire to electrodes embedded in the cochlea in the patient's ear, which in turn stimulate nerves that carry sound to the brain. When cochlear implants became available in 1985, they could help approximately 35 percent of patients; today they can improve hearing in about 80%.

Receiving a cochlear implant, a bionic ear, may seem obviously therapeutic. However, within the deaf community the issue of whether to get a cochlear implant has been controversial. Some deaf individuals have questioned whether these implants are desirable or even therapeutic. Some challenge the standard assumption of the medical community that deafness is a disability that can be "fixed" by having a cochlear implant. At the heart of the debate is the importance and normalcy of the deaf culture. Within the deaf community many find solidarity with others who are deaf and share a common language of signing. Therefore, some in the deaf community believe the acquisition of hearing through cochlear implants needlessly threatens to undermine an adequately functioning culture. These views are held strongly as illustrated by the fact that one member of the deaf community found her tires slashed when she refused to speak out against cochlear implants (Yaffe, 1999).

Some in the deaf community believe that deafness is a disability, but they maintain that it is not worth correcting given the damaged caused to the deaf community. But another position is that deafness is not a disability at all given the availability and success of sign language within the community. Much of the debate about cochlear implants turns on whether one takes the absence of hearing as a disability. If it is, then having a cochlear implant is therapeutic, and, if it is not, then having a cochlear implant is an enhancement.

There are many things that we cannot do, and yet we do not classify them as disabilities. We cannot digest steel and we cannot breathe underwater without special equipment. It would be strange for someone, other than superman, to claim he had a disability because he could not leap tall buildings in a single bound. These sorts of actions are in the realm of inabilities, not disabilities. A disability is a lack of normal ability in reference to a class of individuals. Adults living today are disabled if they do not understand some language, but they are not disabled if they do not understand an extinct language. Those strongly opposed to cochlear implants might argue that within the deaf culture there is normal functioning with full use of language that happens to be a sign language, not an oral one. If the members of the deaf culture are picked as the reference class, then hearing should be regarded as an inability, not a disability. Given this standard, a cochlear implant would not be therapeutic and might be regarded as unnecessary and possibly detrimental.

My purpose here is not to argue for or against cochlear implants. Rather, it is to point out that the distinction between therapy and enhancement is not as straightforward as might be assumed. The decision about getting a cochlear implant is a personal choice and sometimes a difficult one that requires careful consideration of all the consequences. Families of deaf children may find themselves choosing between communities and are sometimes sharply divided within themselves on this issue. The decision can be agonizing because if the implant is to be most effective it must be implanted early in the deaf child's life, preferably before language development occurs.

The cochlear implant debate illustrates that the lack of agreement on what counts as a disability and what does not. However, even if there were agreement on what counts as therapy and what constitutes enhancement, implanted chips can offer a bit of both. For example, an implanted defibrillator can monitor a heart and deliver a shock within 30 seconds after life-threatening irregularities in rhythm are detected. The defibrillator restores normal heart function, but it does so through an enhancing functionality that people without defibrillators do not have. Or, consider an Alzheimer's patient who has a chip embedded that allows her to be located by others and perhaps even guides her back by global positioning satellites. Is this chip therapeutic or enhancing? Suppose a paralyzed patient has a chip implanted that allows him to control the lights in his room by shifting his neural patterns. Is this implant therapeutic or enhancing?

Second Objection: Limitation of Freedom

The second argument against the policy of allowing therapeutic but not enhancing implants is that it arbitrarily limits personal freedom. As long as the implantee and others are not being harmed by the implant, what is the objection to allowing it? In other matters we routinely allow, if not encourage, people to have enhancements. Generally speaking, education enhances as does exercise and a good diet. They enhance the body and the mind and we encourage all of them. Freedom is a core good and we properly allow people the freedom to exercise it. People, at least in a

liberal state, are at liberty to have cosmetic surgery, belly-button rings, and tattoos. Enhancement is what many of us strive for much of the time. As a simple illustration, consider laser eye surgery guided by computer that can enhance vision beyond the normal 20/20. It would seem perverse to insist that a patient should not have the freedom to correct her vision to a better than normal 20/15 but had to stop at 20/20. Similarly, it seems perverse not to give people freedom to enhance themselves in other ways, including the implantation of computer chips if they wish.

In 1998, Kevin Warwick, a cybernetics professor from the University of Reading in the United Kingdom, had a chip implanted that permitted sensors in his laboratory to detect his location and motion. In March, 2002, he had a much more sophisticated computer implant (Warwick, 2002). An array with 100 spikes was implanted in Warwick's wrist to connect his median nerve with a computer. The median nerve travels along the arm and contains both sensory neurons that detect pressure and temperature and motor neurons that connect the spinal cord with muscle groups in the hand. The spikes of the array were implanted in these sensory and motor neurons in the median nerve. Wires from the array traveled up Warwick's arm and surface through a skin puncture in his forearm. The wires were connected to a gauntlet, a transmitting/receiving device, located externally on Warwick's arm. The gauntlet sent information about neural firing to an external computer. The computer, properly calibrated, distinguished neural impulses when Warwick's left hand was open and when it was closed. This provided sufficient binary information for Warwick to guide miniature robots, manipulate a robotic hand, light up specially wired jewelry, and steer an adapted wheelchair.

Warwick could feel the impulse if the information flow was reversed and the computer sent a signal to the gauntlet that transferred it to the implant in his median nerve. In an interesting experiment Warwick wore a baseball cap with an ultrasonic transmitter and receiver. Ultrasonic impulses were sent out from the cap and bounced back quickly if objects were close. In this situation a rapid series of pulses were sent to the computer, to the gauntlet, and to his median nerve. If objects were farther away, the pulses were less rapid. This gave Warwick an extrasensory input. When blindfolded, but hooked up to the ultrasonic device, he could guide himself around his laboratory using bat-like echolocation. In another experiment his wife, Irena, who had a simpler neural connection, and he could exchange binary information back and forth from one nervous system to the other via the Internet. Warwick has raised the possibility that one day more sophisticated information and possibly emotional responses could be communicated from nervous system to nervous system via the Internet.

Warwick's reports on his body image were interesting. Warwick makes it quite clear that having the implant under his skin was important as compared with simply putting on wearable computing that can easily be removed. "[F]rom the very start, I regarded the array and wires as being a part of me. Having it extracted, I knew, would be like losing part of my body, almost an amputation" (Warwick, 2002, p. 292). But he also had some sense of his body's being extended by the machine attachments like the robotic hand. "The articulated hand felt like a part of me, yet, because it was remote, in another sense it didn't" (Warwick, 2002, pp. 233–234).

Warwick acknowledges the potential risks but believes the eventual benefits of computer implant enhancements outweigh these risks. Ethically, should Warwick enhance himself with computer implants? Some believe not. Langdon Winter suggests that such experiments are "profoundly amoral" (Vogel, 2002, p. 1020). Although becoming a cyborg may eventually raise questions about human nature, it is hard to see how the experiments that Warwick performed are beyond straightforward moral judgment. If he is not causing harm to others and not violating any particular duties, why should he not have the freedom to do it? His wife is at some risk of harm, but she freely gave her informed consent. Both his implant procedure and her procedure passed hospital ethics committee evaluation. The experiments may strike some as grotesque, scientifically ill defined, or grandstanding, but such judgments, assuming they were correct, would still not make the experiments amoral or immoral.

THE BORG ARGUMENT

The fear remains that allowing freedom of enhancement through computer chip implants will take us down a slippery slope to some very undesirable results. To imagine a worst case scenario, consider the Borg from the science fiction series *Star Trek*. The Borg is a collection of cyborgs that travels through space in a large cube that has the ability to assimilate new species that it encounters. The Borg's menacing conduct is indicated by its foreboding mottoes: "Resistance is futile" and "We will assimilate you." The inhabitants of the Borg have numerous

unattractive appliances attached to them, have no personal autonomy, and are controlled by the directives o collective consciousness. The inhabitants of the Borg do not have individual lives worth living, at least not intelligent creatures. They neither examine their lives nor personally flourish. And so, the argument runs, we not want to end up like the inhabitants of the Borg. How do we prevent sliding down the slope to such an ex tence once we give people the freedom to implant chips?

Slippery slope arguments are not very convincing, particularly if the slope is rather long and stopping alo it seems possible. There is considerable slope between allowing people the freedom to implant chips and becom ing a Borg culture. But can we easily brake on the slope? I believe we can, but the Borg argument has some forc A Borg culture in which people become slave cyborgs is not something that sane people would choose for them selves. However, other mechanisms might push us toward such a state. Here are two:

The Sleepwalking Scenario: We might inadvertently fall into a Borg-like state if we are not careful. Imagine that people for good reasons decide to have chips implanted in order to communicate with their children or do their jobs better or receive the latest music and sports information or have medication automatically released. Eventually, almost everyone is hooked up to the Web internally and wirelessly. It is the way life is conducted. Babies are given chips as routinely as vaccinations. Such interconnections are useful in organizing our lives. Gradually, for practical, not evil, motives the Web/human system begins to take on a life of its own, coordinating people's activities by sending information tantamount to instructions for where to be and when to be there. Under such a condition the population might look better than the inhabitants of the *Star Trek* Borg, but its behavior might have an uncanny similarity.

The Totalitarian State Scenario: The Borg culture might come into existence through the directives of a dictator of a totalitarian state. Dictators want to control their population. What better way than putting their subjects to work with implants that track their locations and force their labor? Neither the sleepwalking scenario nor the totalitarian state scenario is likely to happen in the immediate future, but these developments are real possibilities. We have yet to produce an Orwellian 1984 society, but that is probably due to a shortage of the right kind of information technology. That technological shortcoming is rapidly being overcome and a Borg culture is something to be on guard against.

Freedom with Responsibility

In general, I am advocating a policy of responsible freedom. People should have the freedom to implant computer chips in themselves, including implants for enhancements. As with all our actions, we should be alert if harm or the risk of harm is a factor. If harm or the risk of harm would occur to either the person being implanted or to others, we need to consider whether the action is justified. Harm does not automatically curtail freedom of action, but it does require justification. Exercising such freedom requires evaluating consequences and formulating relevant policies that can be advocated impartially and publicly so that anyone is permitted to follow them in similar circumstances (Moor, 1999).

Harm can result in many ways when implanting computer chips, but there are three general areas of major concern, three ethical hot spots, to which we should be particularly sensitive in the coming age of cyborgs. These areas are privacy, control, and fairness. Some computer implants will enhance privacy or control or fairness. Some will undermine them.

Privacy

In May 2002, Jeff and Leslie Jacobs and their son, Derek, were the first to have VeriChips implanted in their arms. These chips, little bigger than a grain of rice, store six lines of text. Information is read from the chip by a handheld computer. Such medical information could be lifesaving in giving emergency physicians information about allergies and special medical needs before they administered treatment. In the case of the Jacobs, the chips contain phone numbers and information about previous medications. The U.S. Federal Food and Drug Administration (FDA) ruled a month earlier that it did not regard the chip as a medical device and would not regulate it.

The chip is not very useful unless the implantee is at a hospital that has the appropriate handheld computer reader, but the technology is likely to spread because it is relatively inexpensive. The chip itself is dormant, but

when the right radio frequency energy passes through the skin it activates the chip that in turn emits a radio signal containing an identification number. This number can be sent to an FDA secure data storage site via telephone or the Internet. Given our flourishing Information Age, the demand for implanted chips to store personal, medical, and financial information as well as any information whatsoever is likely to increase.

Implanted chips can be more sophisticated than memory chips. VeriChip is the product of Applied Digital Solutions (ADS) that has for several years been working on another product called the "Digital Angel." The Digital Angel is a tracking device that uses Global Positioning System (GPS) technology. The Digital Angel technology potentially can be used to track almost anyone or anything from children, convicts, and cats to lost hikers and lost luggage. ADS has a bold vision of what the chip might be able to do. According to one early projection, the future chip, if implanted, will be powered by a piezoelectric device that converts energy from normal bodily movements into electricity. The chip will send information to receivers connected to various networks. In addition, the device will be able to collect information about the possessor's body, such as temperature and blood pressure. Blood oxygen and glucose level detection are promised as well. Its designers propose that pulse detection will be based on infrared radiation naturally emitted from the bloodstream. In this vision of the future, solid state accelerometers and gyroscopes will allow the Digital Angel to sense the posture and gait of the possessor to detect sudden falls. EKG and EEG detection are claimed to be in the works. If this information gathering comes to fruition, the objective is to transmit the information to receivers that make it available on the Internet through Web-enabled desktop, laptop, or wireless devices. Depending on the configuration a Digital Angel device could be turned on or off by the possessor, the possessor's body, or remotely by radio signal. The device would not need to be on and transmitting at all times, but would have the ability to turn on automatically if it sensed, for example, a heart attack or was sent an instruction to do so.

Digital Angel is the brainchild of Peter Zhou, who is enthusiastic about the future of implanted chip technology. Despite Zhou's enthusiasm, critics have expressed concerns. Civil rights groups compared the use of implanting these chips to Nazi tattooing and some Christians compared the implanting of the chips to the mark of the beast mentioned in the Bible. Thus far, ADS has brought out the first generation of the Digital Angel as a product to be worn as a wristband or carried. Its initial capability is limited to establishing the location of its possessor.

Regardless of the current stage of development of this implant, the concept of a chip that actively gathers data about its owner as well as sending and receiving information is technologically feasible and such chips will come onto the market for particular uses at some point. Potential uses are plentiful. A person who suffers from arrhythmia could be assisted when the chip monitoring her pulse notifies medical authorities of her location and her condition. A firearm could be programmed to fire only when the chip identifies its user as the proper owner. Herds of animals, not to mention millions of pets, could be tracked so that no animal is lost. Every soldier in a battle unit could be monitored for his or her location and health status. A kidnapped child possessing such a chip could be located and checked for life signs. Such a chip could serve as an ID for business and other human interactions. A potential customer could be positively identified biometrically through transmissions from the chip. Her transaction then would be charged automatically to or deducted from her account based on information passed along by the chip.

A world with implanted personal data chips will generate an enormous flow of personal information in novel ways that will require new protection plans for the privacy of individuals. New policies will need to be created to safeguard the collection of all the up-to-the minute information about people's health, location, financial condition, and other matters transmitted and received by these chips. It is not that personal privacy cannot in principle be protected with the use of such chips. The concern is that the technology will be developed and deployed without establishing privacy protection.

CONTROL

Another ethical hot spot in which implanted chips can provide enormous benefits but put us at risk is control. Respect for the agency of others is a hallmark of ethics. Implanted computer chips hold great promise for both giving and taking away human agency.

In the United States over a million patients suffer from Parkinson's disease, a degenerative neurological disorder that causes them to shake uncontrollably. Another two million suffer from essential tremor that causes

similar violent shaking. The shaking is so debilitating that these patients often have trouble working, eating, and simply getting dressed. In the past the drug L-dopa has been given to Parkinson's patients, though its effectiveness wanes over time. Less than half of the patients with essential tremor are improved with medication. Sometimes patients undergo surgery to destroy parts of their brains that cause the shaking, but this procedure is not reversible and not always effective.

An alternative for these patients is to have a chip implanted. Physicians implant an electrode in a patient's thalamus and run a wire under the scalp to the patient's collarbone where a pulse generator is implanted. This device sends electrical signals customized for each patient to the electrode in the thalamus. A constant stream of electrical shocks blocks the tremors. The device is effective in stopping the shaking in both Parkinson's patients and essential tremor patients. The procedure is reversible and the device can be turned on and off by the patient.

The results of such an implant are nothing short of spectacular. A Parkinson's patient whose hands are shaking violently can run a magnet over his chest activating the pulse generator, and within a few seconds his hands become steady. With another swipe of the magnet, the device is turned off, and his hands will begin to shake again. In one case a typical Parkinson's patient, who had lost her mobility and whose medications made her arms and legs move out of control, could sit down and play complex pieces on the piano after her implant was installed (Freudenheim, 1997).

This Deep Brain Stimulation technique using the pulse generator is now being used for a variety of other medical conditions—even for psychiatric conditions such as obsessive compulsive disorder. One seriously ill patient had repetitive thoughts for hours and would wash his hands seventy times per day. After having a chip implant he stopped his compulsive hand washing and returned to work (Carmichael, 2002).

What is striking about these examples of implants is that they restore agency to the patients. Patients regain control of their lives. The sinister side is the threat that computer implants might be used to remove agency. Chips might be developed to induce uncontrollable shaking, cause obsessive-compulsive disorder.

Consider the recent development of a ratbot. Three electrodes were placed in a rat's brain: two in the somatosensory cortical where the rat's brain processes touch from its right and left whiskers, and one in the medial forebrain bundle where the rat processes pleasure. When one of the two electrodes in the sensory region is stimulated, the rat experiences an apparent touch. If it turns in the direction of its right or left whisker depending on which side is stimulated, it is electrically rewarded in its pleasure center. With this setup and some radio controls to send the signals researchers were able to guide the rat. Using a laptop, researchers maneuvered the rat through a difficult three-dimensional maze that included ladders, filing cabinets, and thin wooden boards. As one researcher appropriately remarked, "I certainly don't think it would be a good idea to put these in primates, or especially in humans" (Cook, 2002).

Rat brains are not human brains, but humans do have pleasure centers in their brains and one can easily imagine the use of implants to control humans. Could such a device be offered to help people stop smoking or lose weight? The military and the penal system might consider using the technology to produce loyal troops and obedient prisoners. Computer implants can potentially elevate human agency or severely reduce it. Continual vigilance regarding the deployment of such devices is necessary to ensure that respect for human agency is maintained.

FAIRNESS

A final ethical hot spot to consider is fairness. Implanted chips can tip the scales of justice in various ways. For example, implanted chips can encourage fairness by giving those with disabilities more power to interact in the world. Consider the case of Johnny Ray who suffered a brainstem stroke in 1997. He has locked-in syndrome and no muscle control. Although he is cognitively intact, he is totally paralyzed and cannot make a motion. Researchers have inserted a subcranial cortical implant. Parts of the implant in the motor cortex are surrounded with tissue culture to encourage brain cells to grow toward the contacts. The patient is asked to think about distinctive conditions such as hot versus cold. The corresponding brain outputs are captured, amplified, and used to control an external device such as a cursor on a computer screen. "By reproducing the same brain pattern, Ray eventually was able to move the cursor at will to choose screen icons, spell, and even generate musical tones" (Hockenberry, 2001, p. 96).

When computer implants improve access and interactive capabilities for those who are disadvantaged, fairness is served. But there are easily imagined situations in which future implants might give the implantee unfair advantages. Just imagine a grandmaster chess chip that contained book openings and generated excellent chess moves. Suppose it were developed, giving its owner superior ability in playing chess. Presumably, such chips would need to be banned in championship play just as steroids are outlawed in Olympic competition. Chip implants that facilitated an athlete's coordination might be banned for similar reasons.

Fairness will be an ongoing concern as chip implants get better and more useful. Eventually, a chip implant divide will emerge between those who have chip implants versus those who do not (Macguire and McGee, 1999). Parents, as parents always do, will want to give their children the best abilities and opportunities possible. Those who can afford chip implants and chip upgrades will have a distinct advantage over those who cannot.

VIEWING THE DISTANT FUTURE

Thus far I have been considering the matter of computer implants in light of common morality in the short run. I have been focusing on ethical concerns for and against implanting chips in the near future. Now I wish to reverse direction and consider the possible implications of computer implants on metaphysical issues and ethical theory in the long run.

The possibility of enhancing humans through computer implants raises the question of what human nature could and should be. Traditionally, essentialist philosophers like Aristotle maintain that humans have a fixed nature. Some existentialist philosophers like Sartre argue that existence precedes essence and that human nature is radically free. We can change our essence by making different choices. In an era of increasing understanding of genetics and neurology, neither position seems quite right. Human nature does not appear to be irrevocably fixed or completely open. Computer implants offer us an opportunity to adjust at least some of our nature. Our essence as humans may not be radically open, but, if we are clever enough in developing implants, we can, if we choose, significantly change our nature from what it is now.

Accurate prediction of what computer implants will be available in the distant future is, of course, impossible. But let's speculate a bit. With implants we can change our internal functioning in ways that are not possible, using variations of our genetic code. We might enhance our sense of sight to access parts of the electromagnetic spectrum far beyond what any humans or other animals can. Similarly our sense of hearing could be radically extended. Artificial devices for touch, taste, and smell already exist and these senses could be great enhanced. We could develop new senses. We might continue to experiment with implants for echolocation, for example, to discover at least in part what it is like to be a bat.

Communicating with other humans may be more direct than ever before. We could have sensors installed in our bodies that would let us know if our loved ones were in danger. We could lock and unlock doors, turn appliances on and off, and adjust the heat in our houses through computers that monitor neural patterns. And our education could be greatly enhanced with better memory and more accurate recall (Eisenberg, 2002). Perhaps education and physical conditioning could be done by downloads, not tedious schooling and training. Improvements our abilities to create music or make inferences might be possible. Although all of this is speculation, nothing seems to preclude these possibilities. However the future develops, it seems likely that human nature as we can readily conceive it could be modified.

Martha Nussbaum, defending an Aristotelean position, has argued against such aspirations.

> . What my argument urges us to reject as incoherent is the aspiration to leave behind altogether the constitutive conditions of our humanity, and to seek for a life that is really the life of another sort of thing—as if it were higher and better life for *us*. It asks us to bound our aspirations by recalling that there are some very general conditions of human existence that are also necessary conditions for the lives that we know, love, and appropriately pursue. (Nussbaum, 1990, p. 379)

Of course, everything depends on what constitutes the conditions of our humanity. The example she uses to make the point is [illegible]'s decision to choose the life of a mortal human being who [illegible] immortality and agelessness, a life of