been put away. He was carving from the images in his mind. His eyes and hands knew where every line, curve, mass must emerge, and at what depth in the heart of the stone to create the low relief." (*The Agony and the Ecstasy.* Doubleday, 1961: 6, 144.)

5. First Foresight Conference on Nanotechnology in October 1989, a talk titled "The Future of Computation." Published in Crandall, B. C. and James Lewis, editors. *Nanotechnology: Research and Perspectives.* MIT Press, 1992: 269. See also *www.foresight.org/Conferences/MNT01/Nano1.html.*
6. In his 1963 novel *Cat's Cradle,* Kurt Vonnegut imagined a gray-goo-like accident where a form of ice called ice-nine, which becomes solid at a much higher temperature, freezes the oceans.
7. Kauffman, Stuart. "Self-replication: Even Peptides Do It." *Nature,* 382, August 8, 1996: 496. See *www.santafe.edu/sfi/People/kauffman/sak-peptides.html.*
8. Else, Jon. *The Day After Trinity: J. Robert Oppenheimer and the Atomic Bomb* (available at *www.pyramiddirect.com).*
9. This estimate is in Leslie's book *The End of the World: The Science and Ethics of Human Extinction,* where he notes that the probability of extinction is substantially higher if we accept Brandon Carter's Doomsday Argument, which is, briefly, that "we ought to have some reluctance to believe that we are very exceptionally early, for instance in the earliest 0.001 percent, among all humans who will ever have lived. This would be some reason for thinking that humankind will not survive for many more centuries, let alone colonize the galaxy. Carter's doomsday argument doesn't generate any risk estimates just by itself. It is an argument for *revising* the estimates which we generate when we consider various possible dangers." (Routledge, 1996: 1, 3, 145.)
10. Clarke, Arthur C. "Presidents, Experts, and Asteroids." *Science,* June 5, 1998. Reprinted as "Science and Society" in *Greetings, Carbon-Based Bipeds! Collected Essays, 1934–1998.* St. Martin's Press, 1999: 526.
11. And, as David Forrest suggests in his paper "Regulating Nanotechnology Development," available at *www.foresight.org/NanoRev/Forrest1989.html,* "If we used strict liability as an alternative to regulation it would be impossible for any developer to internalize the cost of the risk (destruction of the biosphere), so theoretically the activity of developing nanotechnology should never be undertaken." Forrest's analysis leaves us with only government regulation to protect us—not a comforting thought.
12. Meselson, Matthew. "The Problem of Biological Weapons." Presentation to the 1,818th Stated Meeting of the American Academy of Arts and Sciences, January 13, 1999. (*minerva.amacad.org/archive/bulletin4.htm*)
13. Doty, Paul. "The Forgotten Menace: Nuclear Weapons Stockpiles Still Represent the Biggest Threat to Civilization." *Nature,* 402, December 9, 1999: 583.
14. See also Hans Bethe's 1997 letter to President Clinton, at *www.fas.org/bethecr.htm.*
15. Hamilton, Edith. *The Greek Way.* W. W. Norton & Co., 1942: 35.

2.3.4 Nanoscience, Nanotechnology, and Ethics: Promise and Peril

RAY KURZWEIL

Inventor of the first music synthesizer in 1982, entrepreneur, author, and futurist Ray Kurzweil is one of the leading thinkers about contemporary and future technology and its impact on society. In this reading, he argues that the pace of technological change is exponential rather than linear, as many people implicitly believe, and that the twenty-first century will witness a major technological revolution springing from the convergence of the science of genetics, nanotechnology, and robotics (GNR) and artificial intelligence. However, like all new technologies, GNR is a "double-edged sword" that holds great promise and great peril for human civilization. One worry, expressed by Bill Joy (Selection 2.3.3) is that self-replicating nanobots will escape into the environment and cause severe and irreparable damage to the natural world. Although he takes such threats seriously, Kurzweil believes that the development and introduction of GNR is inevitable and that broad relinquishment of these technologies is not feasible, but that it is also possible to develop ethical guidelines for "fine-grained" relinquishment of certain kinds of particularly dangerous GNR applications. While continuing to be optimistic about the promise of these technologies, he believes that it is possible for society to control them so as to enjoy their benefits while avoiding the dangers they pose.

FOCUS QUESTIONS

1. What are the three features of technological evolution that lead Kurzweil to conclude that the rate of technological change is exponential rather than linear?
2. Kurzweil suggests that the introduction of some kinds of nanotechnology devices is inevitable. What reasons does he adduce that support this conclusion?
3. At one point, Kurzweil refers to the views of Ted Kaczynski (aka the Unabomber) and, perhaps surprisingly, seems to endorse some aspects of his view of technology. What part of Kaczynski's view does he share? What part does he disagree with?
4. Why is Kurzweil opposed to the relinquishment of broad areas of technology? What kinds of "fine-grained" relinquishment does he think are feasible?
5. Do you agree with Kurzweil that it is possible for society to enjoy the benefits of twenty-first century GNR technologies while mitigating and controlling the risks? Explain.

KEYWORDS

defensive technologies GNR, Moore's law, nanobots, nanotechnology, paradigm shift, relinquishment

Our rapidly growing scientific and technological ability to manipulate matter and energy at ever smaller scales promises to transform virtually every sector of society, a phenomenon that presents manifest ethical responsibilities. There will be increasing overlap between nanotechnology and other technologies, such as biotechnology and artificial intelligence. And as with these previous scientific and technological transformations, we will be faced with deeply intertwined promise and peril.

THE NANO-FRONTIER

Nanoscience and nanotechnology today have been expanded to include essentially any science or technology where the key features are measured in a modest number of nanometers (under 100 by some definitions). By this standard, contemporary electronics has already passed this threshold. Eric Drexler has further developed the concept of building molecule-scale devices using molecular assemblers that would precisely guide chemical reactions by means of information. Moreover, just as technologies related to information develop at an exponential pace, generally doubling in capability and price-performance every year, so the size of technology is itself inexorably shrinking, and most of technology will be "nanotechnology" by the 2020s.

This era will bring us the ability to essentially convert software, that is, information, directly into physical products. We will be able to

Source: *Encyclopedia of Science, Technology, and Ethics,* Vol. A–C, ed. Carl Mitcham. Macmillan Reference. New York: Thomson Gale, 2006, xli–xlvii.

produce virtually any product for pennies per pound. Computers will have greater computational capacity than the human brain, and we will be completing the reverse engineering of the human brain to reveal the software design of human intelligence. We are already placing devices with narrow intelligence in our bodies for diagnostic and therapeutic purposes. With the advent of nanotechnology, we will be able to keep our bodies and brains in a healthy, optimal state indefinitely. Nanotechnology and related advanced technologies will bring us the opportunity to overcome age-old problems, including pollution, poverty, disease, and aging.

Many object to the intermingling of the so-called natural world with the products of our technology. However, the increasing intimacy of our human lives with our technology is not a new story. Human life expectancy was thirty-seven years in 1800. Most humans at that time lived lives dominated by poverty, intense labor, disease, and misfortune. We are immeasurably better off as a result of technology, but there is still a lot of suffering in the world to overcome. We have a moral imperative, therefore, to continue the pursuit of knowledge and of advanced technologies that can continue to overcome human affliction. There is also an economic imperative to continue.

Nanotechnology is advancing on hundreds of fronts. We cannot relinquish its pursuit without essentially relinquishing all of technology, which would require acts of totalitarianism inconsistent with the values of our society. Technology has always been a double-edged sword, and that is certainly true of nanotechnology. However, we will have no choice but to confront the challenge of guiding nanotechnology in a constructive direction. Any broad attempt to relinquish nanotechnology will only push it underground, which would interfere with the benefits while actually making the dangers worse.

With the human genome project, three to five percent of the budgets were devoted to the ethical, legal, and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive. Near-term applications of nanotechnology are more limited in their benefits and more benign in their potential dangers. We cannot say a priori that all nanoengineered particles are safe, nor would it be appropriate to deem them necessarily unsafe. Environmental tests thus far have not shown reasons for undue concern.

I believe that existing regulatory mechanisms are sufficient to handle near-term applications of nanotechnology. As for the long term, we need to appreciate that a myriad of nanoscale technologies are inevitable. The current examinations and dialogues on achieving the promise while ameliorating the peril are appropriate and will deserve increased attention as we get closer to realizing these revolutionary technologies.

THE NANO-BACKGROUND: MODELS OF TECHNOLOGY TRENDS

Models of technology trends show that nanotechnology and related advanced technologies are inevitable. They are deeply integrated into our society and are advancing on many diverse fronts, comprised of hundreds of small steps, each benign in itself.

Intuitive Linear and Historical Exponential Views Although exponential trends did exist a thousand years ago, they were at that very early stage where it is so flat and so slow that it looks like no trend at all. Today, everyone expects continuous technological progress and the social repercussions that follow. But the future will nonetheless be far more surprising than most observers realize because few have internalized the fact that the rate of change itself is accelerating.

Most long-range forecasts of technical feasibility underestimate the power of future developments because they are based on the "intuitive linear" view of history rather than the "historical exponential" view. We will not experience a hundred years of progress in the twenty-first century; rather we will witness on the order of twenty thousand years of progress (at today's rate of progress). An unexamined intuition provides the impression that progress changes at the rate that we have recently experienced because an exponential curve approximates a straight line when viewed for a brief duration.

But an assessment of the history of technology shows that technological change is exponential. Indeed, we find "double" exponential growth, meaning that the rate of exponential growth is itself growing exponentially. These observations are based on a rich model of diverse technological processes.

The Law of Accelerating Returns The ongoing acceleration of technology is the inevitable result of the "law of accelerating returns," which describes the acceleration of the pace and the exponential growth of the products of an evolutionary process, including technology, particularly information technologies.

The law of accelerating returns has three key features. First, evolution applies positive feedback as the more capable methods resulting from one stage of evolutionary progress are used to create the next stage. As a result, the rate progress of an evolutionary process increases exponentially over time, as the "returns" of that process (e.g., speed or cost-effectiveness) increase exponentially. As an evolutionary process becomes more effective, greater resources are invested in it, resulting in a second level of exponential growth (i.e., the rate of exponential growth itself grows exponentially).

A second feature is "technological paradigm shifts." A specific paradigm (a method or approach to solving a problem) provides exponential growth until the method exhausts its potential. When this happens, a paradigm shift (a fundamental change in the approach) occurs, which enables exponential growth to continue. Each paradigm follows an "S-curve," which consists of slow growth, followed by rapid growth, followed by a leveling off as the particular paradigm matures. During this third phase in the life cycle of a paradigm, pressure builds for the next paradigm shift. The acceleration of the overall evolutionary process proceeds as a sequence of S-curves, and the overall exponential growth consists of this cascade of S-curves.

A third key feature is that the resources underlying the exponential growth of an evolutionary process are relatively unbounded. One resource is the order of the evolutionary process itself. Each stage of evolution provides more powerful tools for the next. The other required resource is the "chaos" of the environment in which the evolutionary process takes place and which provides the options for further diversity. In technological evolution, human ingenuity and the ever-changing market sustain innovation.

The evolution of life forms and technologies constantly accelerates. With the advent of a technology-creating species, the exponential pace became too fast for evolution through DNA-guided protein synthesis and moved on to human-created technology. Technology goes beyond mere tool making; it is a process of creating ever more powerful technology using the tools from the previous round of innovation. The first technological steps took tens of thousands of years. For people living in this era, there was little noticeable technological change. By 1000 C.E., progress was much faster and a paradigm shift required only a century or two. The nineteenth century saw more technological change than in the nine centuries preceding it. Then in the first twenty years of the twentieth century, we saw more advancement than in all of the nineteenth century. Now, paradigm shifts occur in only a few years. The paradigm shift rate is currently doubling every decade. So the twenty-first century will see about a thousand times greater technological change than its predecessor.

Moore's Law and Beyond The exponential trend that has gained the greatest public recognition has become known as "Moore's Law." Gordon Moore, one of the inventors of integrated circuits, noted in the mid-1970s that we could squeeze twice as many transistors on an integrated circuit every twenty-four months. Given that the electrons have less distance to travel, the circuits also run twice as fast, providing an overall quadrupling of computational power.

However, the exponential growth of computing is much broader than Moore's Law. If we plot the speed per price of forty-nine famous calculators and computers spanning the twentieth century, we note that there were four paradigms that provided exponential growth in the price-performance of computing before integrated circuits. Therefore, Moore's Law was the fifth paradigm to exponentially grow the power of computation. When Moore's Law reaches the end of its S-curve, the exponential growth will continue with three-dimensional molecular computing, constituting the sixth paradigm.

Moore's Law narrowly refers to the number of transistors on an integrated circuit of fixed size. But the most appropriate measure to track is computational speed per unit cost. This takes into account many levels of innovation in computer design. For example, there are many nascent technologies that build circuitry in three dimensions in a way that mimics the parallel organization of the human brain. One cubic inch of nanotube circuitry would be a million times more powerful than the human brain. There are more than enough new computing technologies now being researched to sustain the law of accelerating returns as applied to computation.

Specific paradigms do ultimately reach levels at which exponential growth is no longer feasible. That is why Moore's Law is an S-curve. But the growth of computation will continue exponentially. Paradigm shift, or innovation, turns the S-curve of any specific paradigm into a continuing exponential. A new paradigm takes over when the old paradigm approaches its natural limit.

Other Technologies There are many examples of the exponential growth implied by the law of accelerating returns in technologies as varied as DNA sequencing, communication speeds, brain scanning, electronics of all kinds, and even in the rapidly shrinking size of technology. Exponential growth in communications technology has been even more explosive than in computation. Miniaturization is a trend that will have profound implications for the twenty-first century. The salient implementation sizes of technologies, both electronic and mechanical, are shrinking at a double-exponential rate.

The future nanotechnology age will result not from the exponential explosion of computation alone, but rather from the synergies that will result from intertwined technological revolutions. Every point on the exponential growth curves represents an intense human drama of innovation and competition. It is remarkable that these chaotic processes result in such smooth and predictable exponential trends.

EXAMPLES OF TRUE NANOSCIENCE AND NANOTECHNOLOGY

Ubiquitous nanoscience and nanotechnology is two to three decades away. One forthcoming achievement will be "nanobots," small robots the size of human blood cells that can travel inside the human bloodstream. There have already been successful animal experiments using this concept.

In addition to human brain reverse engineering, these nanobots will be able to perform a broad variety of diagnostic and therapeutic functions inside the human body. Robert Freitas, for example, has designed robotic replacements for human blood cells that perform thousands of times more effectively than their biological counterparts. His "respirocytes" (robotic red blood cells) could allow one to sprint for fifteen minutes without taking a breath. His robotic macrophages will be far more effective than our white blood cells at combating pathogens. His DNA repair robot would be able to repair DNA transcription errors, and even implement needed DNA changes. Although Freitas' conceptual designs are two or three decades away, there has already been progress on bloodstream-based devices.

Nanobot technology has profound military applications, and any expectation that such uses will be relinquished is highly unrealistic. Already, the U.S. Department of Defense (DOD) is developing "smart dust," or tiny robots to be used for surveillance. Billions of invisible spies could monitor every square inch of enemy territory and carry out missions to destroy enemy targets. The only way for an enemy to counteract such a force is with their own nanotechnology. Nanotechnology-based weapons will obsolete weapons of larger size.

In addition, nanobots will be able to expand our experiences and our capabilities. Nanobot technology will provide fully immersive virtual reality by taking up positions in close proximity to every interneuronal connection related to the senses. If we want to enter virtual reality, the nanobots suppress all of the inputs coming from the real senses, and replace them with the signals that would be appropriate for the virtual environment.

Scientists at the Max Planck Institute have developed "neuron transistors" that can detect the firing of a nearby neuron, or alternatively, can cause a nearby neuron to fire, or suppress it from firing. This amounts to two-way communication between neurons and the electronic-based neuron transistors. The scientists demonstrated their invention by controlling the movement of a living leech from their computer.

The Internet will provide many virtual environments to explore. We will be able to "go" to these virtual environments and meet others there, both real and simulated people. Of course, ultimately there will not be a clear distinction between the two. By 2030, going to a web site will mean entering a full-immersion virtual-reality environment, encompassing all of the senses and triggering the neurological correlates of emotions and sexual experiences.

"Experience beamers" circa 2030 will beam a person's entire flow of sensory experiences and emotions. We'll be able to go to a web site and experience other people's lives. Full-immersion visual-auditory environments will be available by 2010, with images written directly onto our retinas by our eyeglasses and contact lenses. The electronics will be embedded in our glasses and woven into our clothing, so computers as distinct objects will disappear.

The most significant implication of nanotechnology and related advanced technologies of the twenty-first century will be the merger of biological and nonbiological intelligence. Nonbiological intelligence is growing at a double-exponential rate and will vastly exceed biological intelligence well before the middle of this century. However, in my view, this nonbiological intelligence should still be considered human, as it is fully derivative of the human-machine civilization.

Our brains are relatively fixed in design, but brain implants based on massively distributed intelligent nanobots will ultimately expand our memories a trillion fold and improve all of our cognitive abilities. Since the nanobots are communicating with each other over a wireless network, they can create any set of new neural connections, break existing connections, create new hybrid biological-nonbiological networks, and add new nonbiological networks.

Using nanobots as brain extenders is a significant improvement over surgically installed neural implants. Nanobots will be introduced without surgery and can be directed to leave, so the process is easily reversible. They can change their configuration and alter their software. Perhaps most importantly, they are massively distributed and can take up billions or trillions of positions throughout the brain, whereas a surgically introduced neural implant can only be placed in a few locations.

THE ECONOMIC IMPERATIVES OF THE LAW OF ACCELERATING RETURNS

The economic imperative of a competitive marketplace is driving science and technology forward and fueling the law of accelerating returns, which, in turn, is transforming economic relationships. We are moving toward nanoscale, more intelligent machines as the result of many small advances, each with their own particular economic justification.

There is a vital economic imperative to create smaller and more intelligent technology. Machines that can more precisely carry out their missions have enormous value. There are tens of thousands of projects that are advancing the various aspects of the law of accelerating returns in diverse incremental ways. Regardless of near-term business cycles, the support for "high tech" in the business community has grown enormously. We would have to repeal capitalism and every visage of economic competition to stop this progression.

The economy has been growing exponentially throughout this century. Even the Great Depression of the 1930s represented only a minor blip compared to the underlying pattern of growth. Recessions, including the Depression, represent only temporary deviations from the underlying curve. Statistics in fact greatly understate productivity growth (economic output per worker), which has also been exponential.

Inflationary factors are offset by the double-exponential trends in the price-performance of all information-based technologies, which deeply affect all industries. We are also undergoing massive disintermediation in the channels of distribution through the Internet and other new communication technologies and escalating efficiencies in operations and administration. Current economic policy is based on outdated theories that do not adequately model the size of technology, bandwidth, megabytes, intellectual property, knowledge, and other increasingly vital constituents that are driving the economy.

Cycles of recession will not disappear immediately. However, the rapid dissemination of information, sophisticated forms of online procurement, and increasingly transparent markets in all industries have diminished the impact of these cycles. The underlying long-term growth rate will continue at a double-exponential rate. The rate of paradigm shift is not noticeably affected by the minor deviations caused by economic cycles. The overall growth of the economy reflects completely new forms of wealth and value that did not previously exist: nanoparticle-based materials, genetic information, intellectual property, communication portals, web sites, bandwidth, software, data bases, and many other new technology-based categories.

Another implication of the law of accelerating returns is exponential growth in human knowledge, including intellectual property, education, and learning. Over the course of the long twentieth century we increased investment in K-12 education by a factor of ten. We have a one hundred fold increase in the number of college students. Automation has been eliminating jobs at the bottom of the skill ladder while creating new and better paying jobs at the top. So, the ladder has been moving up, and we have been exponentially increasing investments in education at all levels.

PROMISE AND PERIL

Science and technology have always been double-edged swords, bringing us longer and healthier life spans, freedom from physical and mental drudgery, and many new creative possibilities, while at the same time introducing new and salient dangers. We will need to adopt strategies to encourage the benefits while ameliorating the risks. Relinquishing broad areas of technology, as some critics have proposed, is not feasible, and attempts to do so will only drive technology development underground, which will exacerbate the dangers.

As technology accelerates toward the full realization of biotechnology, nanotechnology and "strong" AI (artificial intelligence at or above human levels), we will see the same intertwined potentials: a feast of creativity resulting from greater human intelligence combined with many new dangers. Nanobot technology requires billions or trillions of such intelligent devices to be useful. The most cost-effective way to scale up to such levels is through self-replication. A defect in the mechanism curtailing nanobot self-replication could be disastrous. There are steps available now to mitigate this risk, but we cannot have complete assurance in any strategy that we devise today.

Other primary concerns include "Who is controlling the nanobots?" and "Who are the nanobots talking to?" Organizations or individuals could put undetectable nanobots in water or food supplies. These "spies" could monitor and even control thoughts and actions. Existing nanobots could be influenced through software viruses and other software "hacking" techniques. My own expectation is that the creative and constructive applications of this technology will dominate, as they do today. But we need to invest more heavily in developing specific defensive technologies.

There are usually three stages in examining the impact of future technology: awe at its potential to overcome problems; then a sense of dread at a new set of dangers; followed by the realization that the only viable and responsible path is to set a careful course that can realize the promise while managing the peril.

Bill Joy, cofounder of Sun Microsystems, has warned of the impending dangers from the emergence of self-replicating technologies in the fields of genetics, nanotechnology, and robotics, or "GNR." His concerns include genetically altered designer pathogens, self-replicating entities created through nanotechnology, and robots whose intelligence will rival and ultimately exceed our own. Who's to say we will be able to count on such robots to remain friendly to humans? Although I am often cast as the technology optimist who counters Joy's pessimism, I do share his concerns regarding self-replicating technologies. Many people have interpreted Joy's article as an advocacy of broad relinquishment, not of all technology, but of the "dangerous ones" like nanotechnology. Joy, who is now working as a venture capitalist with the legendary silicon valley firm of Kleiner, Perkins, Caufield & Byers investing in technologies such as nanotechnology applied to renewable energy and other natural resources, says that broad relinquishment is a misinterpretation of his position and was never his intent. He has recently said that the emphasis should be to "limit development of the technologies that are too dangerous," not on complete prohibition. He suggests, for example, a prohibition against self-replicating nanotechnology, which is similar to the guidelines advocated by the Foresight Institute.

Others, such as Bill McKibben, the environmentalist who was one of the first to warn against global warming, have advocated relinquishment of broad areas such as biotechnology and nanotechnology, or even of all technology. However, relinquishing broad fields would be impossible to achieve without essentially relinquishing all technical development.

There are real dangers associated with new self-replicating technologies. But technological advances, such as antibiotics and improved sanitation, have freed us from the prevalence of such plagues in the past. We may romanticize the past, but until fairly recently, most of humanity lived extremely fragile lives. Many people still live in this precarious way, which is one reason to continue technological progress and the economic enhancement that accompanies it. Should we tell the millions of people afflicted with devastating conditions that we are canceling the development of all bioengineered treatments because there is a risk that these same technologies may someday be used for malevolent purposes? Most people would agree that such broad-based relinquishment is not the answer.

The Relinquishment Issue Relinquishment at the right level is part of a responsible and constructive response to these genuine perils. The issue, however, is: At what level are we to relinquish technology? Ted Kaczynski (the Unabomber) would have us renounce all of it. This is neither desirable nor feasible. McKibben takes the position that many people now have enough wealth and technological capability and should not pursue more. This ignores the suffering that remains in the human world, which continued technological progress could alleviate.

Another level would be to forego certain fields (such as nanotechnology) that might be regarded as too dangerous. But such sweeping strokes of relinquishment are untenable. Nanotechnology is the inevitable result of the persistent trend toward miniaturization that pervades all of technology. It is not a single centralized effort, but is being pursued by a myriad of projects with many goals.

Kaczynski argued that modern industrial society cannot be reformed because technology is a unified system in which all parts are dependent on one another. It is not possible to get rid of the "bad" parts of technology and retain only the "good" parts. He cited modern medicine as an example, arguing that progress depends on several scientific fields and advancements in high-tech equipment. Kaczynski was correct on the deeply entangled nature of the benefits and risks, but his overall assessment of the relative balance between the two was way off. Joy and I both believe that technology will and should progress, and that we need to be actively concerned with the dark side. Our dialogue concerns the granularity of relinquishment that is feasible and desirable. Abandonment of broad areas of technology will only push them underground where development would continue unimpeded by ethics and regulation. In such a situation, it would be the less-stable, less-responsible practitioners who would have all the expertise.

One example of relinquishment at the right level is the proposed ethical guideline by the Foresight Institute that nanotechnologists agree to relinquish the development of physical entities that can self-replicate in a natural environment. Another is a ban on self-replicating physical entities that contain their own codes for self-replication. Such entities should be designed to obtain codes from a centralized secure server, which would guard against undesirable replication. This "broadcast architecture" is impossible in the biological world, which represents one way in which nanotechnology can be made safer than biotechnology. Such "fine-grained" relinquishment should be linked to professional ethical guidelines, oversight by regulatory bodies, the development of technology-specific "immune" responses, as well as computer assisted surveillance by law enforcement agencies. Balancing privacy rights with security will be one of many challenges raised by some new nanotechnologies.

Computer viruses serve as a reassuring test case in our ability to regulate nonbiological self-replication. At first, concerns were voiced that as they became more sophisticated, software pathogens had the potential to destroy computer networks. Yet the "immune system" that has evolved in response to this challenge has been largely effective. Although self-replicating software entities do cause damage from time to time, no one would suggest we do away with computers and the Internet because of software viruses. This success is in a highly productive industry in which there is no regulation, and no certification for practitioners.

Defensive Technologies and The Impact of Regulation. Arguments such as McKibben's for relinquishment have been influential because they paint a picture of future dangers as if they were released into an unprepared world. But the sophistication and power of our defensive technologies and knowledge will grow along with the dangers. When we have "gray goo" (unrestrained nanobot replication), we will also have "blue goo" ("police nanobots"). We cannot say with assurance that we will successfully avoid all misuse. We have been able to largely control harmful software virus replication because the requisite knowledge is widely available to responsible practitioners. Attempts to restrict this knowledge would have created a far less stable situation.

The present challenge is self-replicating biotechnology. By reprogramming the information processes that lead to and encourage disease and aging, we will have the ability to overcome these afflictions. However, the same knowledge can also empower a terrorist to create a bioengineered pathogen.

Unlike biotechnology, the software industry is almost completely unregulated. Although bioterrorists do not need to put their "innovations" through the FDA, scientists developing defensive technologies are required to follow regulations that slow innovation. It is impossible under existing regulations and ethical standards to test defenses to bioterrorist agents on humans. Animal models and simulations will be necessary in lieu of infeasible human trials, but we will need to go beyond these steps to accelerate the development of defensive technologies.

We need to create ethical and legal standards and defensive technologies. It is quite clearly a race. In the software field the defensive technologies have remained ahead of the offensive ones. With extensive regulation in the medical field slowing down innovation, this may not happen with biotechnology.

There is a legitimate need to make biomedical research as safe as possible, but our balancing of risks is skewed. The millions of people who need biotechnology advances seem to carry little political weight against a few well-publicized casualties from the inevitable risks of progress. This equation will become even starker with the emerging dangers of bioengineered pathogens. We need a change in public attitude in terms of tolerance for necessary risk.

Hastening defensive technologies is vital to our security. We need to streamline regulatory procedures to achieve this. However, we also need to greatly increase our investment explicitly in defensive technologies. In the biotechnology field, this means the rapid development of antiviral medications.

The comparable situation will exist for nanotechnology once replication of nano-engineered entities has been achieved. We will soon need to invest in defensive technologies, including the creation of a nanotechnology-based immune system. Such an immune system may itself become a danger, but no one would argue that humans would be better off without an immune system because of the possibility of autoimmune diseases. The development of a technological immune system for nanotechnology will happen even without explicit efforts to create one.

It is premature to develop specific defensive nanotechnologies as long as we have only a general idea of the threat. However, there is a dialogue on this issue, and expanded investment in these efforts should be encouraged. The Foresight Institute, for example, has devised a set of ethical standards and strategies for assuring the development of safe nanotechnology. They are likely to be effective with regard to preventing accidental release of dangerous self-replicating nanotechnology entities. But the intentional design and release of such entities is more challenging.

CONCLUSION

Protection is not impossible, but we need to realize that any level of protection will only work to a certain level of sophistication. We will need to continue to advance the defensive technologies and keep them ahead of the destructive technologies. The challenge of self-replication in nanotechnology impels us to continue the type of study that the Foresight Institute has initiated. With the human genome project, three to five percent of the budget was devoted to the ethical, legal and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive. Science and technology will remain double-edged swords, and the story of the twenty-first century has not yet been written. We have no choice but to work hard to apply these quickening technologies to advance our human values, despite what often appears to be a lack of consensus on what those values should be.

BIBLIOGRAPHY

Drexler, Erik K. (1986). *Engines of Creation*. Garden City, NY: Doubleday.

Drexler, Erik K. (1992). *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. New York: Wiley.

Freitas, Robert, Jr. (1999). *Nanomedicine*. Vol. 1: *Basic Capabilities*. Georgetown, TX: Landes Bioscience.

Joy, Bill. (2000). "Why the Future Doesn't Need Us." *Wired* 262.

Kurzweil, Ray. (1999). *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Viking.

Kurzweil, Ray. (2005). *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

McKibben, Bill. (2003). *Enough: Staying Human in an Engineered Age*. New York: Henry Holt.

2.4 BIOTECHNOLOGY AND GENETIC ENGINEERING

2.4.1 A Glimpse of Things to Come

LEE M. SILVER

Princeton University molecular biologist Lee Silver opened his 1997 book *Remaking Eden: Cloning and Beyond in a Brave New World* with this provocative discussion of what human reproduction might look like in 2010, 2050, and 2350. He predicts that reproductive genetic engineering technology will make it possible for lesbian couples to have children that are genetically related to both "parents," that genetic resistance to diseases such as AIDS will be woven into an embryo's DNA, that human cloning will become widely accepted, and that, in the further future, the human race will divide into two classes: the Naturals and the