4

SMALL IS POWERFUL

Daniel Sarewitz and Edward Woodhouse

The metal stirrup, which migrated from Asia to western Europe in the eighth century, allowed the energy of a galloping horse to be directly transmitted to the weapon held by the man in the saddle—a combat innovation of devastating impact. In those days, horses and tack were costly, possessed almost exclusively by landowners. Battlefield prowess and wealth thus went hand in hand; together they fostered the traditions of a "warrior aristocracy" and laid the foundations for European feudal society itself. When the Anglo-Saxon king Harold prepared to defend Britain against invading Normans in 1066, he actually dispensed with his horse and ornamental wooden stirrups, choosing to lead his numerically superior forces on foot. The outnumbered Normans, however, boasted a stirrup-equipped cavalry, and thus won the day—and the millennium.[1]

Such narrative has the ring of mythology, yet the experience of the industrialized world reinforces the idea that, when innovators deploy new tools to their advantage, they change society in the process. The invention of the cotton gin in the late eighteenth century allowed a vast expansion of cotton cultivation in the American South—directly fueling a resurgence in the importation of slaves for plantation labor. One hundred and fifty years later, the mechanical cotton picker suddenly rendered obsolete the jobs of millions of African American sharecroppers, and catalyzed a thirty-year migration of 5 million people out of the rural South and into the cities of the North. Given emerging engineering expertise, the mechanical cotton picker may have been inevitable, but it proliferated rapidly because plantation owners feared the civil rights movement and welcomed a technological replacement for the exploitation labor upon which they depended.[2]

Technology and society thus evolve together. The stirrup emerged in tandem with feudalism, agricultural equipment cannot be understood apart from the legacy of slavery and labor issues, and nuclear weapons joined with U.S. and Soviet hegemonies to constitute a prime determinant of geopolitical evolution after World War II. Cars, television, air conditioning, and birth control likewise arose in particular social contexts and contributed to the remaking of everyday life.

If innovations as apparently modest as the metal stirrup and the cotton gin can transform society to its roots in a period of decades or less, what of technologies now on the horizon that aim to revolutionize the very processes by which new materials are designed and produced, that are blurring the boundary between the inanimate and the living, that may combine the powers of machine intelligence and human consciousness? No one can fully understand the long-term implications of such advances, emerging under the heading of nanotechnology—"the art and science of building complex, practical devices with atomic precision," with components measured in nanometers, billionths of a meter.[3] But rapidly improving capacities for miniaturization are now combining with continuing refinements in computation, mechatronics, and telecommunications. Innovations based on nanotechnology may interact to rival the combined epoch-making social effects of chemicals, nuclear missiles, mechanized transport, computerized data processing, antibiotics, TV, and agribusiness. Society a century from now might be barely recognizable—and the range of possibilities goes from wonderful to dire.

The essence of the nanotechnology story is the continuation of a fifty-year trend of machine miniaturization culminating in the rise of design control at the molecular level. Nanotechnology is not confined to a single area of innovation; "smallness" is its unifying attribute. Researchers in a number of technical fields are keenly interested in manipulation of matter at the nanoscale, and funding is assured because many of the research forefronts hold promise for business and military applications.

As one small measure of this growing interest, U.S. government support for nanotechnology research increased sixfold between 1997 and 2003, to $710 million per year. Because most of the research is at a precommercial stage, much of this funding aims at hastening the journey from research to application. For example, the U.S. National Science Foundation (NSF) supports university nanotechnology research centers to explore the fundamental science and engineering that is supposed to enable rapid innovation. What is the promise of these investments?

- As conventional silicon chips approach their maximum capacity for memory storage, nanotechnology offers the potential of chiplike functionality from single molecules. Scientists foresee very small, inexpensive computers with thousands of times more computing capacity than current machines, perhaps introducing a second computer revolution that could dwarf the changes of the past fifty years.[4]
- Advances in nanofabrication are leading directly to a new generation of sensors, with "surfaces that can sense and bind to chemical and biological agents [and] emit a measurable electrical or optical signal when binding occurs." This could lead to a "reliable, inexpensive, and portable way to ensure that the world's drinking water and food supplies are free from contaminants"; sensors in homes and workplaces "that could detect minute quantities of all biological and chemical hazards and provide appropriate safety measures if detected"; and devices "as small as the tip of a hypodermic needle" that "could detect thousands of diseases."[5]
- Scientists are working "to evolve organisms to live with and work with other kinds of inorganic materials. . . . The project is working with viruses that can be engineered to stick to various elements. . . . The viruses can grow in sheets, creating a flexible surface holding nanoparticles of various materials. . . . This could lead to flexible computer displays,

> while removing the viruses after a nanostructure is formed could expand its usage into conditions where biological materials fail."[6] Other researchers are seeking to replicate biological functions with synthetic ones, designing and synthesizing organic molecules and supramolecular arrays that can mimic green plants' photosynthetic processes—perhaps opening up a new age powered by solar energy in a far more fundamental sense than what that term means at present.[7]

IBM and Xerox are among an increasing number of large corporations engaged in nanotechnology R&D, and start-up firms hoping to mimic the explosive success of Silicon Valley are racing to get products onto the market. Carbon Nanotechnologies, for example, claims to be a "world leading producer of single-wall carbon nanotubes . . . the stiffest, strongest, and toughest fibers known." Their most advanced product, Bucky-Plus™ Fluorinated Single-wall Carbon Nanotubes, goes for $900 per gram, more than fifty times the price of gold.[8] Nanomix is working to develop "new hydrogen storage systems that will power the fuel cell revolution, by using nanostructured materials to store solid-state hydrogen for automotive and portable power applications."[9]

To the nanotechnology research community and its advocates, the future looks bright indeed. As one well-known technological visionary, Newt Gingrich—chair of the NanoBusiness Alliance and former speaker of the U.S. House of Representatives—puts it:

> Nanotechnology, the science of developing tools and machines as small as one molecule, will have as big an impact on our lives as transistors and chips did in the past forty years. Imagine highly specialized machines you ingest, systems for security smaller than a piece of dust and collectively intelligent household appliances and cars. The implications for defense, public safety and health are astounding.[10]

Even normally staid government reports burst with promotional fervor—"Forward-looking researchers believe they could end up with synthetic creations with lifelike behaviors"[11]—apparently relegating those who suffer discomfort at such prospects to the ranks of the "backward-looking." To the visionaries of nanotechnology, "[o]ur world is riddled with flaws and limitations. Metals that rust. Plastics that break. Semiconductors that can't conduct any faster. And so on. Nanotechnology can make it all better—literally—by re-engineering the fundamental building blocks of matter. It is one of the most exciting research areas on the planet, and it may lead to the greatest advances of this century."[12]

Yet when a National Science and Technology Council report says, "If present trends in nanoscience and nanotechnology continue, most aspects of everyday life are subject to change,"[13] what exactly are the authors thinking? "Science discovers, technology creates, man adapts," as the 1939 New York World's Fair proclaimed? The use of the passive voice is wonderfully revealing: the world *is* to be transformed by inevitable, autonomous, disembodied processes called science and technology, but no one, apparently, is doing the transforming.

Not exactly. Thus far, the exuberant decision to remake the world with nanotechnology has come from committees drawn from a small group of experts, mostly male, mostly upper middle class, mostly North American, universally in possession of great technical expertise. But twenty-first-century nanoscientists and engineers have thought no more carefully about the social aspects of their work than had the previous century's technologists who introduced nuclear weaponry and nuclear reactors to the world, or the chemists who blithely synthesized millions of tons of chlorinated chemicals without regard for their ecological and health effects.[14] Still lacking is a recognition that evolving sensible paths of advance, paths that can win wide public support and ensure broad public benefit, requires time for patient deliberation. The intelligence of democracy is sustained by debate and negotiation among partisans with partially conflicting values, different competencies, and different institutional bases and interests. Experts alone cannot supply these diverse perspectives.[15]

At this point, much work in nanotechnology is no more than a reflection of the joy that scientists and engineers experience when they use new tools to do new things of interest to them:

> Donald Eigler of IBM's Alden Research Center remembers the day in 1990 when he and Erhard K. Schweitzer, who was visiting from the Fritz-Haber Institute in Berlin, moved individual atoms for the first time. In his laboratory notebook Eigler used big letters and an exclamation mark to write "THIS IS FUN!" Using one of the most precise measuring and manipulating tools the world had ever seen, the researchers slowly finessed thirty-five xenon atoms to spell out the three-letter IBM logo atop a crystal of nickel. To be sure, it only worked in a vacuum chamber kept at a temperature that makes the North Pole seem tropical.[16]

Fun indeed. "Nanotechnology researchers love their newfound ability to move atoms on surfaces."[17] The enthusiasm is understandable, but there is something disquieting about the promise that the joy of doing science will translate into a world inevitably transformed for the better. Delight in tinkering is thus revealed as the justification and foundation for ushering in social changes of unknown kinds and potentially unlimited extent. Beneath the surface, moreover, lies a political reality: to continue to do the work that gives them such personal and professional rewards, scientists and the agencies that support them may simply be saying what they think elected officials want to hear, in order to enhance prospects for future funding.[18] Upon such banal motives is the world remade.

Technological revolutions do not build the world from scratch, of course. Products of innovation are introduced into society through institutions and systems that already exist, and whose strengths and flaws are likely to persist even in the face of rapid technological change. Consider, for example, the health care arena. Despite continual promises to the contrary, experience shows that new biomedical technologies tend to offer benefits only in exchange for higher overall costs, while contributing to well-documented inequities of health care access, delivery, and outcomes, at least in the United States.[19]

Nanoscale techniques are predicted to "revolutionize the speed with which new compounds are screened for therapeutic potential as new drugs. . . . If the trend is similar to that of microelectronics, the rate could grow exponentially."[20] Pharmaceutical companies' R&D trajectories, however, are oriented primarily toward generating revenue, not toward improving health. For example, since the mid-1970s, of the more than 1,200 drugs put on the market by the pharmaceutical industry, only four were aimed at infectious diseases of the tropics, such as malaria, that kill millions annually. Of these four drugs, moreover, two originally were developed for other purposes, and one has since been withdrawn from the market.[21] We can expect a nanotech-enabled proliferation of new drugs that can help people in affluent societies cope with everything from neurosis to impotence to the asymptotic decline of our aging bodies, but unless present motivations for science and innovation change, we should anticipate little benefit for those whose needs are greatest.

Nanotechnology also promises to accelerate the trend toward diagnosis of diseases where there is no cure. "With arrays of ultra miniaturized sensors that sample a range of chemicals or conditions, the confidence level and specificity of detection would be much greater than is now possible with separate macroscopic sensors."[22] For example, DNA sensors will soon have the capability of screening for multiple diseases, "including sexually transmitted diseases, cystic fibrosis, and genetic predisposition to colon cancer and blood hypercoagulation."[23] Cystic fibrosis remains incurable, of course, and knowing about predispositions to cancer has already been revealed, in the case of breast cancer, to be at best a mixed blessing.

As nanosensors begin to detect the first molecular indicators of a disease, moreover, they will certainly help save lives, but they will also lead to increased medical interventions that are unnecessary or actively harmful. The ongoing controversies over the effectiveness of mammography for breast cancer and prostate-specific antigen (PSA) tests for prostate cancer provide a window into what may happen on a larger scale. While both tests provide early detection capability, statistical evidence from numerous clinical trials has not demonstrated that the new techniques actually save or extend lives.[24] They do, however, trigger the demand for additional tests and treatments, some of which are debilitating in their own right.

Nanosensors promise to move detection to even earlier stages—at the first molecular manifestation of elevated PSA levels, for example—despite the fact that most prostate tumors grow so slowly that they are not life threatening.[25] More generally, abnormal genetic and chemical processes occur in the normal body all the time—and normal body processes are often sufficient to take care of them before they become serious. Nanosensing of the earliest harbingers of diseases thus promises to stimulate a significant rise in unnecessary treatments and significant side effects. At the same time, the advent of nanodetection capabilities will considerably expand the information that insurance providers can use in making decisions about coverage. As competitive businesses, insurance companies often try to increase their profitability by denying coverage to those at high risk, a goal that governments can only partially thwart via regulation (except by providing publicly funded, universal medical coverage).

Of course, some applications of medical nanotechnology may be worth the costs and risks, while others will not be. The question is how to distinguish, how to act on the distinctions, and who should be involved in the selection process. The same need for making careful choices will arise in assessing the promises of nanotech for other sectors: How can innovations in computing help alleviate, rather than exacerbate, the widespread experience of information overload in modern life? How can manufacturing be restructured to enhance the quality of work life—and avoid the marginalization of millions of workers worldwide, which industrial innovation has often caused in the past? Can nanotech be used to increase the intelligence and autonomy of military hardware (thereby keeping soldiers off the front lines) without encouraging less well-equipped enemies to turn toward unprotected civilian targets like the World Trade Center?

Into the Unknown

The hype surrounding medical diagnostic applications of nanotech may be naively (or cynically) optimistic, but the consequences are likely to lie in the realm of the familiar. If, however, nanotechnology achieves its ultimate potential—to literally assemble materials and machines on a molecule-by-molecule basis, and to achieve functionality at the level of individual molecules—then we will be moving into a realm in which we have no experience. Indeed, "nanofabrication" is one of the chief areas of emphasis for government-funded nanotech research. Progress in this realm may ultimately lead to what is known as an "assembler," a device "that is able to manufacture almost anything. . . . Fed with simple chemical stocks, this amazing machine breaks down molecules, and then reassembles them into any product you ask for."[26] Right now this is obviously the stuff of science fiction, and some knowledgeable observers are confident it will remain so. Others disagree, however, and the Zyvex Corporation, which claims to be "the first molecular nanotechnology company," is pursuing this holy grail of nanotechnology, a system "capable of manufacturing bulk materials or arbitrary structures with atomic precision, getting nearly every atom in the desired place."[27]

Probably the first coherent warning about nanotechnology of this new kind came from nanoscientist and technology forecaster Eric Drexler, whose 1986 book *Engines of Creation,* while for the most part extolling the prospects of the technology, devoted one chapter to possible dangers:

> Powered by fuels or sunlight, [replicating assemblers] will be able to make almost anything (including more of themselves) from common materials. . . . [A]ssembler-based replicators could beat the most advanced modern organisms. "Plants" with "leaves" no more efficient than today's solar cells could out-compete real plants, crowding the biosphere with an inedible foliage. Tough, omnivorous "bacteria" could outcompete real bacteria: they could spread like blowing pollen, replicate swiftly, and reduce the biosphere to dust in a matter of days.[28]

Drexler's work did not gain much public attention, but a similar version of the hypothetical dangers of nanotechnology made front pages in 2000 when Sun Microsystems chief scientist Bill Joy published in *Wired* magazine an article titled "Why the Future Doesn't Need Us." Joy described a world of self-replicating, exponentially proliferating "nanobots" that could drown the planet in an uncontrollable "gray goo." Because of his standing as one of the leading architects of the world's information infrastructure, his warning made waves: Joy was no Luddite. While initial media reports of his pessimistic view were respectful, the research and technology communities quickly mobilized like antibodies to neutralize him. Nobel Prize winner Richard Smalley said: "My advice is, don't worry about self-replication nanobots. . . . It's not real now and will never be in the future."[29] John Armstrong, retired IBM vice president for research, swatted Joy away without even bothering to name him: "If you are worried, as some seem to be, about a robotic future full of nano mechanisms that don't need us, I suggest you rent a copy of Woody Allen's *Sleeper* from the video store, and restore your sense of balance!"[30]

Adherents of what sometimes verges on a new nanotechnology religion seem not to notice their own intemperance. The point, surely, is not whether Joy's specific worries come true; it is that his predictions are not obviously less reasonable extrapolations of current trends than those made by nanotech promoters. Joy's concerns arguably deserve special consideration, moreover, because they identify a downside that outweighs any reasonable estimate of the new technology's positive potential. Nor does Joy have an intrinsic conflict of interest, as do many of the researchers who stand to benefit if their promotional activities translate into research funding.

Less controversial than Joy's predictions about gray goo is his recognition that nanotechnology will fuel a second revolution in computer power that could lead to a hybridization of human and machine intelligence. While some technological visionaries view this as the desirable and inevitable result of human invention, other people may not feel entirely sanguine about launching or accelerating technological changes that could make humans as we now think of them . . . well, obsolete. *Does* the future need us? Quantum computers and human brains may combine to create something entirely new. One articulate champion of this vision is the inventor Ray Kurzweil:

> We are entering a new era. I call it "the Singularity." It's a merger between human intelligence and machine intelligence that is going to create something bigger than itself. It's the cutting edge of evolution on our planet. One can make a strong case that it's actually the cutting edge of the evolution of intelligence in general, because there's no indication that it's occurred anywhere else. To me that is what human civilization is all about. It is part of our destiny and part of the destiny of evolution to continue to progress ever faster, and to grow the power of intelligence exponentially. To contemplate stopping that—to think human beings are fine the way they are—is a misplaced fond remembrance of what human beings used to be. What human beings are is a species that has undergone a cultural and technological evolution, and it's the nature of evolution that it accelerates, and that its powers grow exponentially, and that's what we're talking about. The next stage of this will be to amplify our own intellectual powers with the results of our technology.[31]

Kurzweil's enthusiastic portrayal of exponential growth of machine intelligence betrays a peculiar understanding of what matters in the world, "what human civilization is all about": the continued evolution of intelligence. Is humanness really so tied up with ever-increasing information-processing ability? It is easy to imagine a species with more powerful brains than ours—science fiction authors do it all the time—but whether or not they are "human" is another matter. One need only read Homer or Shakespeare to recognize that the essence of humanity has, for better or worse, survived the industrial and information revolutions pretty much intact. The past generation or two of exponential growth of information-processing capabilities so far doesn't appear to have made political, economic, and technological elites discernibly smarter about how they wield their newfound powers; that kind of wisdom is not derived from analytical prowess. (In fact, an excess of confidence in the power of rational analysis has underlain such disasters of modernity as central planning in the Soviet Union, U.S. involvement in the Vietnam War, and the replacement of natural forests with monocultures.)

Whether "the Singularity" advances the cause of human well-being or retards it will reflect not the technologies themselves but the regimes under which they are wielded. Nanotechnology advocates appear just about oblivious to this simple truth. We can be fairly confident that predictions of nanotech-enabled future utopias (and dystopias) will someday seem as quaint (or as malign) as More and Verne and Marx seem to us now.

Steering Lessons

The naiveté dominating nanotechnology discussions fits with one of the best-researched conclusions from historians of technology: One ought never accept experts' rosy predictions about any emerging technological potential. As the political scientist Langdon Winner puts it, contemporary technoscientists tend to work within the "mastery tradition" in Western thought and practice—embracing an assumption that knowledge can and will be used to conquer "nature."[32] This mechanistic view of the universe, a legacy of the seventeenth-century optimism that gave birth both to modern science and to modern democracy, is evident in pronouncements made on behalf of nanotechnology (but with democracy nowhere to be seen). Although temporarily sobered by the environmental surprises and nuclear near-catastrophes of the twentieth century, most scientific and technical fields appear governed by what psychoanalysts would refer to as denial and overcompensation—reiterating ever more loudly the mantra that technical progress leads reliably to social progress.[33] The mastery tradition, in other words, is alive and well despite the bruising experiences and partially successful social movements of the past hundred years.

Kurzweil and other nanotechnology visionaries give the mastery tradition a new twist. Technological evolution, they believe, is largely autonomous, proceeding on paths that can be little altered by human choice. Thus, nanotechnology commentaries all generally share the contradictory idea that specific technological changes are coming inexorably, and yet people are going to be freer than ever, better than ever. This is an incredible scenario: an automatic and inevitable translation of technological destiny into an improved life for everyone.

The proclamations made on behalf of nanotechnology are rooted in a thermodynamic and philosophical absurdity: that control at the micro level translates into control at the macro level. Indeed, the word "control" is central to the promise of nanotech: "Nanotechnology's relevance is underlined by the importance of *controlling* matter at the nanoscale for healthcare, the environment, sustainability, and almost every industry."[34] The real world—the experienced world, the world in which humans must make decisions about, say, how to make use of nanotechnology—is made up of complex systems comprising innumerable components interacting in ways that are often intrinsically unpredictable. This is another of the most firmly established understandings developed by systematic social science analysis of technological innovation: unintended consequences often are greater than those foreseen and intended by innovators.

Unintended consequences emerge in part because control exercised at one level very often leads to unpredictable reactions on another level. Nanotechnology is the ultimate application of reductionism, and its power to confer control over increasingly small components of nature may prove great indeed. But just as, for example, the meso-scale control afforded by the automobile and the coal-burning power plant yield macro-level consequences that include air pollution, climate change, and the geopolitics of fossil fuels, so we can be reasonably sure that specific nanotechnology applications will have impacts not readily controlled or even understood by those creating or using them. As with transforming technologies ranging from the stirrup to the production line to the cotton gin to the hydrogen bomb, we ought to expect that the unanticipated consequences, both good and bad, will provoke profound social disruption.

If, as promised, nanotechnology revolutionizes our whole system of manufacturing and invention, for example, then the impacts will be enormous and unpredictable. To get a sense of possible scale, consider only one aspect of the first industrial revolution—the transformation of human labor. Prior to the nineteenth century, even the most economically advanced societies were predominantly agrarian and rural. For the majority of people, work was rooted in the home and the family. Vagaries of weather and transportation imposed hardship, but most people and families possessed a diversity of skills that gave them autonomy from wage labor and resilience to cope with a variety of challenges. In hard times, resort to subsistence farming and barter was usually possible.[35]

The advent of modern manufacturing machinery changed that. Increasing industrialization and urbanization linked workers to both labor and consumer markets, removing their need and ability to maintain the diverse skills that had been prerequisites for survival in the preindustrial world. Labor itself became a commodity, subject to the same fluctuations and influences as other commodities. During an economic downturn, factories fired people or closed down entirely; for the first time, workers could not easily respond to changing economic conditions by switching to a different type of work or moving to a subsistence mode. As political economist Karl Polanyi observed, "To separate labor from other activities of life and to subject it to the laws of the market was to annihilate all organic forms of existence and to replace them by a different type of organization, an atomistic and individualistic one."[36] Some observers then and now considered this to be wonderful, others found it to be terrible; what is beyond dispute is that the sociotechnical changes exerted profound influence on the pace, character, and meaning of millions of lives.

The glib claim that humanity will gain greater control over the material world via nanotechnologies thus ignores our historical knowledge that technology changes, in unforeseeable ways, the social contexts within which humans can act. Moreover, "humanity" is not an actor, and hence cannot control anything. The point is that certain individuals and their organizations make choices about certain limited domains, and new technical capacities ordinarily give some people and organizations enhanced means of exercising influence—both over the material world and over other people. So the real question is not whether nanotechnology will allow humanity to assert better control over nature, but *which* humans will be making the choices, how well their values and aims will match those of the majority, whether they will respect minority needs, and whether they will adopt strategies to protect against the sorts of problems and dislocations we have here discussed—and those we cannot yet imagine.

Sometime in the next few years a human being probably will be cloned. This will occur despite widespread repugnance for such activities, including legal proscriptions in the United States on federally funded human cloning. The technical momentum appears unstoppable, however, and the political will to outlaw human cloning outright seems to be lacking. Given the spread of the technologies to such fringe groups as the Raelians, a quasi-religious organization with members around the world, a ban would be impossible to enforce. The time to prevent human cloning would have been about 1980, when the emerging biotechnology revolution began to point toward the capability—but before it existed, so no one would have had much to lose from a ban.

Yet at that time there would have been little stomach among scientists for governmental regulations aimed at forestalling particular lines of research on the basis of public sentiment, and those advocating early action would undoubtedly have been countered by accusations of "antiscientism" and "Luddism." And that is precisely our point: when major innovations loom, neither government institutions, nor scientific practice, nor consumer markets are set up to act sensibly in the public interest in a timely way. Here is how Langdon Winner framed the question in *The Whale and the Reactor:*

> In an age in which the inexhaustible power of scientific technology makes all things possible, it remains to be seen where we will draw the line, where we will be able to say, here are possibilities that wisdom suggests we avoid. I am convinced that any philosophy of technology worth its salt must eventually ask, How can we limit modern technology

to match our best sense of who we are and the kind of world we would like to build?[37]

What, then, would it take to do better?

Wise steering of any technology requires coping with two issues that are largely absent from the promotional juggernaut behind nanotechnology: disagreement and uncertainty. Over the past half-century, political scientists studying public decision-making have concluded that expertise is no substitute for political negotiation because analysis alone can never settle questions as complex as those involved in directing technological innovation. These scholars also have found that, for similar reasons of complexity and unforeseeable consequences, trial-and-error learning is the best approach for introducing major new technologies into society. But how the negotiations are approached makes a great deal of difference, as do the strategies and tactics by which innovators and regulators address the process of learning from experience.

The first step toward more intelligent trial-and-error governance of nanotechnology is to bring it clearly onto the public agenda, and this seems to be under way.[38] In 2002, Michael Crichton published a nonfiction essay on nanotechnology in *Parade* magazine, together with his novel *Prey,* which sensationalized and brought under the national spotlight scenarios such as those advanced by Drexler and Joy.[39] A new Center for Responsible Nanotechnology started operation in December. One of our most thoughtful observers of science and technology, Princeton physicist Freeman Dyson, acknowledged the potentially "grave dangers" of nanotechnology in an article in *The New York Review of Books,*[40] and at about the same time the University of Toronto's Joint Centre for Bioethics warned of a nanotechnology backlash. Nanotechnology is creeping into the public eye.[41]

The virtues of greater openness ought, by now, to be obvious to science and technology planners. Regardless of one's opinions about the merits and risks of nuclear power and genetically modified foods, for example, it is hard to imagine that those involved in the early promotion of these technologies would not now wish that they had more aggressively engaged the public in developing a vision for what was wanted and how it should be pursued.

Yet familiar pathologies are already starting to play themselves out in nanotechnology. While media stories are beginning to stimulate interest in nanotech, several public interest groups, including the Science and Environmental Health Network and the ETC Group, have begun to voice concerns, and have been criticized for doing so. At present, there is no real forum for interested parties to discuss their perspectives. Congressional hearings on the subject have mostly been sycophantic exercises dominated by insider experts; they could have been titled "Hoorah for Nanotechnology!" As the experience with nuclear energy revealed, few governments readily provide a forum for dissidents. The result is greater polarization of rhetoric on all sides, and a hardening of positions over time that makes deliberative action increasingly difficult.

Neither an automobile nor a conversation nor an emerging technology can be steered properly if it is moving too fast for those nominally in charge to learn and adjust on the basis of feedback. To facilitate the broad public deliberation that will be necessary for wise pursuit and deployment of nanotechnology, there appears to be no reasonable alternative to slowing down certain aspects of research and commercialization. This may seem a radical and unprecedented idea, but it actually is neither. When biotechnology research began in the 1970s, some of the key scientists involved declared a moratorium on what was then known as "recombinant DNA" research, and then gathered at the Asilomar conference center in California to work out precautionary guidelines. The National Institutes of Health subsequently endorsed these guidelines, which became the de facto regulatory framework for research in that arena. And although the Asilomar process was flawed because it was expert-driven and applied only to the immediate risks of laboratory research—neglecting consideration of the long-term social risks of the products of research—it at least suggested that the scientific community could engage in a process of responsible self-assessment.

Given the huge uncertainties about the future social impacts of nanotechnology, we ought to think of the unfolding revolution as a grand experiment—a clinical trial—that technologists are conducting on society. From this perspective, we can reflect upon the robust societal consensus that demands prior informed consent as a basis for participation in scientific experiments. This consensus is formally codified in the World Medical Association's Helsinki Declaration, strengthened most recently in 2000, and reinforced in the public consciousness by the memory of, for example, the Tuskegee experiments, where African American males with syphilis were left untreated as part of a "control group," despite the existence of treatments known to be efficacious.[42]

In the United States, every publicly funded research project involving human subjects is monitored by an institutional review board (IRB) that must approve the research before it can be conducted. Every university, independent laboratory, and private-sector lab receiving federal funding for human subjects research has an IRB; there are thousands of boards operating in the United States, nearly 800 in California alone. These boards demonstrate that comprehensive governance is a reasonable goal, and while IRBs certainly impose a cost in terms of the efficiency of conducting research, they are an accepted element of a scientific infrastructure that respects human dignity. Similar commitments of the entire research enterprise to larger democratic strictures occur in experiments with animals and in compliance with environmental health and safety regulations. Comprehensiveness, in other words, is possible, when the stakes are high and societal intent is clear.

But human subjects research is more than an illustrative example. The implications of the nanotechnology revolution dwarf those of any particular clinical trial or psychological test—yet we accept this experiment on society with no moral compunction, no mechanism of oversight, no obligation to understand what we are doing while we are doing it. IRBs provide a model that could be expanded into a broad social assessment and consent mechanism attached to all major nanotechnology programs, especially those likely to introduce entirely novel processes and products into society. Such an approach could bring a variety of social-impact-assessment tools, such as scenario-building and technology foresight, together with well-accepted deliberative processes such as consensus conferences and town meetings, to create a comprehensive but decentralized approach to public participation in technology steering.

Assessment and public discussion cannot substitute for regulation, however. One possible regulatory model is the Premanufacture Notification system for new chemicals. In 1976, after terrible experiences with vinyl chloride, PCBs, DDT, and other chlorinated chemicals, Congress passed the Toxic Substances Control Act mandating that all new chemicals be approved by the Environmental Protection Agency prior to manufacture and use. The Food and Drug Administration, of course, has a similar system for pharmaceuticals that is even more elaborate and restrictive. The rationale for both regimes is that the public cannot rely on technologists and manufacturers alone to judge the safety of new products, that this is a job for government regulators who will not be biased by anticipated profits. Although the private sector will reactively oppose any such regulatory scheme, we note that the most highly regulated industries, such as chemicals and pharmaceuticals, are also among the most innovative, profitable, and competitive.

As products are approved for manufacture and distribution, nanotechnology regulators need to learn from the doleful experience of nuclear power and not scale up too quickly. Nuclear technology foundered in large part because too many reactors were built too fast, before utilities, government, and society in general could learn from experience with pilot projects and smaller reactors. Most utilities that owned reactors found out too late that the costs were far higher than anticipated, the machinery far more finicky, and the regulatory environment much more troublesome than alternative energy-generation methods would have been in that era. Going slowly while learning about a young technology is exactly the opposite of what market buyers and sellers tend to do, of course, so there is simply no substitute for tight governmental regulation to achieve this goal.

The growing nanotechnology controversy also underscores the stupidity of having dissolved the congressional Office of Technology Assessment in 1995. It simply is not possible to govern wisely a technological civilization without intermediary agencies mandated to draw together the best thinking of a wide range of relevant experts, stakeholders, and interest groups. The only such organizations in the United States today (for instance, the National Research Council) are too beholden to their scientific constituencies—which include, of course, those working on nanotechnology—to act as honest brokers in the process.

In sum, perhaps what best describes what we are after is a quality of "reflexiveness," a process by which the broadened community of participants concerned about the direction and impacts of scientific advance and technological innovation gain a fuller understanding of the social context within which they operate. This new understanding necessarily becomes an improved basis for making decisions about how and where to move forward. Yet the unpredictability of nanotechnology-enabled futures means that we will need more than broadened self-awareness. We will also need to assess the emerging implications and impacts of nanotechnology in real time—as new principles, products, and processes are developed but before they proliferate—and apply the results of what we have learned to our deliberative fora, regulatory structures, and research institutions. This sort of intelligent trial-and-error process can allow us to learn from experience at an acceptable cost. It is an amazing irony of technological innovation that for all our rational intelligence applied in the short term, we trust our long-term well-being to the hope that the good unanticipated consequences of our inventions will outweigh the bad. As the potency of our technologies continues to accelerate, this seems more than imprudent.

Finally, then, consider Bill Joy's one unassailable argument: miniaturization coupled with increased power for computing and biomanipulation makes opportunities for doing mischief more available, less expensive, and less dependent on complex institutional research infrastructure than ever before. Until recently, weapons of mass destruction required huge laboratories costing billions of dollars. At the other end of the scale was Ted Kaczynski, the Unabomber, working out of a shack in the wilderness and proceeding victim by victim with technologies that had been available for generations, diffused through the conventional mail. These two extremes are converging. The 2001 anthrax mailer was probably operating alone, or in a very small group, out of a modest facility. Shift ahead to the coming era of Bill Joy's thought experiment: instead of a disgruntled mathematician, imagine that the Kaczynski of the twenty-first century is a disgruntled molecular biologist who creates a designer pathogen with no antidotes, and diffuses it throughout the population via nanoscale drones. At this point, the probability of such an occurrence is impossible to estimate, but given the potential consequences, wouldn't a wise civilization do its best to reduce that probability to zero?