Chapter 11
Don't Count Society Out

John Seely Brown & Paul Duguid

John Seely Brown is Chief Scientist and former director of the Xerox Palo Alto Research Center. He has co-authored several major best-selling books with Paul Duguid, the most recent of which is "The Social Life of Information."

Overview

We are going to discuss alarming, disarming, and arming society against high-technology threats.

As Bill Joy has argued, new technologies pose a profound challenge to society, particularly in their potential for self-replication. It is vital to recognize the problems posed by such things as bioengineering, nanotechnology, and robotics. But, we argue here, to prevent possible threats looming out of proportion and solutions, by contrast, shrinking to relative insignificance, it is also important to set them in a suitable context. The pervasive tunnel vision that surrounds most technological prophecy, however, makes it very hard to see much context at all. Consequently, the technologically driven road to the future can seem unavoidable. Noting the fallibility of this sort of in-the-tunnel prediction, we suggest that the environment in which technology is developed and deployed is too often missing from discussions of both promise and problems. Yet, it is to this environment that society must turn to prevent the annihilation that Joy fears. We suggest that society should look to the material environments on which any new, self-replicating form of life must depend to impose limits on that replication. But we also suggest that it must consider the social environment in which technologies emerge. Transformative new technologies may call for new institutions, often transformative in themselves, and it is as important for society to apply itself as consciously and collectively to the development of the latter, as it does to the development of the former.

Revolutions

Whatever happened to the household nuclear powerpack? In 1940, Dr. R.M. Langer, a physicist at Cal. Tech, unhesitatingly predicted that domestic atomic power would introduce a technological and social revolution "*in our own time*." Uranium-235, he assured his readers, would provide heat, light, and energy from a household supply.[1] A nuclear plant the size of a typewriter would power cars—though these would have little use, as family nuclear rockets would be better for most journeys.[2] Completing the revolutionary picture, journalists, economists, and even the president looked forward to a time when "atoms for peace" would provide "zero cost" fuel to generate electricity, allowing rich and poor countries equal access to the benefits of cheap energy.[3] Consequently, many believed with Langer that U-235 would produce a society without class, privilege, cities, nations, or geographical boundaries whatsoever. Instead, there would be a "single, uniformly spread community, sharing as neighbors the whole surface of the world."

For the digital age, the technology has changed. Many of the utopian predictions remain eerily the same. Indeed, from the telegraph to the Internet, pundits have predicted that technological innovation would drive social change.[4] Of course, there is a certain historical specificity to Langer's predictions. The cataclysmic destruction of Hiroshima and Nagasaki brought the idea of a nuclear world to most people's attention. In their aftermath, forecasters attempted to erase the horror and allay society's deepest doubts with visions of a utopian future.[5] Nonetheless, the nuclear example illustrates a couple of general points about such predictions. First, it indicates the fallibility of techno-social predictions (and yet, paradoxically, their curious predictability—worldwide community seems forever to be the promise of the latest gizmo). Second, it reveals the two standard voices that discuss technology's effects on society—one wildly optimistic, one dourly pessimistic. As dramatically new technologies appear, two opposed choruses quickly form, each intoning their part in Dickens famous antiphonal song of revolution:

> It was the best of times. It was the worst of times.

It's easy—and usual—to cast the second voices as Luddites. Yet the nuclear debate makes it clear that things are not so simple. One of the earliest voices of doubt about nuclear technology, after all, was Albert Einstein, who called for "watchfulness, and if necessary quick action" on the part of government in his famous letter to President Roosevelt in 1939. And one of the strongest voices after the war was J.R. Oppenheimer. From his unique position of unrivalled scientific knowledge and personal experience, and with a remarkable ethical sensibility, Oppenheimer warned of the dangers inseparable from a nuclear age. (Though he was Langer's colleague at Cal Tech, Oppenheimer was not impressed by such predictions: "We do not think," he wrote bluntly in 1945, "automobiles and airplanes will be run by nuclear units."[6]) Einstein and Oppenheimer. These are hardly Luddites.

Oppenheimer was invoked both directly and indirectly by Bill Joy in his anti-millenarian article "Why the Future Doesn't Need Us" in the April, 2000 issue of *Wired*. Joy, too, is clearly no Luddite. (Luddites may have gleefully jumped on the Joy bandwagon—but that's a different matter.) Among other things, he played essential roles in development of Berkeley UNIX, Sun Microsystems, and the cross-platform Java language.[7] With this experience behind him, he placed himself in his article at the crossroad of the new millennium and asked in essence, Do we know where are we are being driven by modern technological developments? Will the outcomes be good? And if not, Can we change our route or is the direction already beyond our control? Though he refers back to the nuclear debate, Joy insists that his worries are distinct: the nature of the technological threat to society has fundamentally changed. Citing bioengineering, nanotechnology, and robotics as the greatest worries, Joy argues that low-cost, easily manipulated, and ultimately self-replicating technologies present a fundamentally new, profoundly grave, and possibly irreversible challenge to society's continued existence. Consequently, he warns that the "last chance to assert control - the fail-safe point is rapidly approaching." Asserting control, he continues, will be difficult. Not only is the time to act short, but the threat is not well understood, and the old responses and institutions (implicit in Einstein's suggestion of government intervention and control) are inadequate.

No one should doubt the importance of raising these questions. It does remain to ask, as we do here however, whether this debate has been framed in the most useful way—a way that helps us to address the underlying issues? Or has it emerged in a way that the envisioned solutions (such as the individual refusal of scientists of good conscience to work with dangerous technologies) seem weak in comparison to the threat? As we shall argue, the Joy debate looks to the future with penetrating vision, but it is also a form of tunnel vision, excluding in the way it is framed, broader social responses.[8] Such tunnel vision leads to a position where, rather than arming society for the struggle, the debate may not only be alarming society, but unintentionally disarming it with a pervasive sense of near inescapable gloom.[9] Joy describes a technological juggernaut that is leading society off a precipice. While he can see the juggernaut clearly, he can't see any controls: "We are being propelled into this new century," he writes, "with no plan, no control, no brakes." It doesn't follow, however, that the juggernaut is uncontrollable or unstoppable—nor, indeed, that society can only act in frenzied, individualized haste.

In searching for a brake, it's first important to remember, as we did with Langer, the context of the Joy debate, in particular where and when it began. The article appeared, after all, in *Wired*. In tone and substance, an article there is unlikely to resemble Einstein's measured letter to Franklin Roosevelt. For the best part of a decade, *Wired* has been an enjoyable cheerleader for the digerati. During that decade, the digerati, in as much as they are a single group, specialized in almost unchecked technological determinism and euphoria. Whatever was new and digital was good. Whatever was old (except the occasional prognostication of Toffler or McLuhan) was irrelevant. Society was bad, individualism was good. Constraints (and particularly government and institutional constraints) on technology or information were, on first principles, an abomination. The imperative to "embrace dumb power" left no room for hesitation, introspection, and certainly not retrospection. A repeated subtext was implicitly "damn the past and full speed ahead." The inherent logic of technology would and should determine where society is going. Consequently, the shift with Joy's article from cheering to warning marks an important moment in the digital *zeitgeist*. But one for which the digerati and the readers of *Wired* were not well prepared. In a very short period, prognosticators, like investors, came to realize that rapid technological innovation can have a down side. As with many investors, the digerati's savvy tone swept straight from one side in the conventional Dickensian chorus to the other: from wild euphoria to high anxiety or deep despondency—with little pause to look for a middle ground. The results of such a polarized change in mood are not surprising, but equally not necessarily reliable. When they felt that technology was taking society triumphantly forward, the digerati saw little need to look for the brake. Now, when some fear that rather than being carried to heaven, society is being rushed to oblivion, it's not surprising to hear a cry that no brake can be found.

So, in what follows, we try to suggest where we might look for brakes. In reviewing not the Joy debate *per se*, but issues and topics it raises, we argue that brakes and controls on the technological juggernaut probably lie outside the standard narrow causes that drive technological predictions and raise technological fears. In particular, they lie beyond the individuals and individual technologies that form the centerpieces of most discussions of technology and its effects on society. They lie, rather, in that society and, more generally, in the complex environments—social, technological, and natural—in which technologies emerge, on which they ultimately depend, but which, from a technology-centered point of view, can be hard to see. These dense interconnections between humankind, its environment, and its technologies do not guarantee a secure future for humanity by any means. But they offer its best hope for protection and survival.

Tidings of discomfort

Whatever the tendencies of the digerati, there was good reason to raise a clamor over the issues Joy discusses. Despite various points of comparison, his position is significantly different from Oppenheimer's. Much of the work that preceded the nuclear explosions occurred in secrecy or at least obscurity. Consequently, popular perception of the nuclear age began with an unprecedented and terrifying bang. Oppenheimer didn't need to raise anxiety. Rather, as we have seen, most after-the-fact effort went in the other direction, attempting to calm the populace by promising an unceasing flow of good things. The digital age, by contrast, developed the other way around. For all its apparent speed, it has come upon us all rather slowly. Computers had several decades to make their journey from the lab to domestic appliances, and promises of unceasing benefits and a generally good press have cheered them along the way. In 1982, *Time* even made the personal computer its "person" of the year. Similarly, the Internet has had two decades and a good deal of cheerleading of its own to help it spread beyond its initial research communities. Of Joy's worrisome trio, time has also made biotechnology and robotics familiar, and the press they have received—dystopian science-fiction aside—has been predominantly a good one.

Conversely, nanotechnology is so obscure that few have any idea what it is. Consequently, to generate concern—or even discussion—about the issues these technologies raise demands first a good deal of shouting just to get attention.[10]

In this vein, but echoing the digerati's implicit cry that the past is "off," Joy suggests that the challenges we face are unprecedented and the threats almost unimaginable. "We are," he argues,

> On the cusp of the further perfection of extreme evil, an evil whose possibility spreads well beyond that which weapons of mass destruction bequeathed to the nation-states, on to a surprising and terrible empowerment of extreme individuals.

In contrast to bombs, for example, viruses, whether bioengineered or software engineered, pose insidious and invisible threats. And whereas only states had the infrastructure and finances to develop the former, the latter may need only cheap, readily available devices to instigate irreversible threats to all humanity. Able to take advantage of our dense, social and technological interconnections to replicate, these threats will spread from obscure beginnings to devastating ends—as the "I love you virus" spread from student's computer in Manila to cripple computers around the world.

It may nonetheless be something of a disarming overstatement to suggest that these positions are unprecedented. Different forms of biological warfare, from poisons to germs, have been with us for a long time and, as we shall see, can teach us something about future threats. Moreover, even the optimistic Dr Langer saw there would be problems if Uranium-235 got into the hands of "eccentrics and criminals."[11] Importantly, Joy is more worried about these things getting not so much into malicious hands, as out of hand altogether and beyond the reach of people. Here the underlying fear concerns self-replication, the ability of bioengineered organisms, nano devices, or robots to reproduce geometrically, without (and even despite) the intervention of humans, and at unimaginable speeds. This fear recalls the nuclear scientists' similar worries about unstoppable, self-sustaining chain reactions that might turn our relatively inert planet into a self-consuming star.[12] The most frightening cases, then, involve individual technologies that, once released, neither need nor respond to further human intervention, whether malicious or well intentioned, but on their own can replicate themselves to such an extent that they threaten human existence. How are we to deal with these? Joy wants to know.

And how should we reply?

Digital endism, technological determinism

First, we need to set such fears in context of digital-technology predictions more generally, of which these worries form a part. The fear that "the future doesn't need us" is the ultimate example of what we have called "endism." This is the tendency of futurists in general and the digerati in particular to insist that new technologies will bring an absolute and irrevocable end to old ways of doing things.[13] Business consultants, for example, have told their clients to "forget all they know" and redesign their operations in entirely new ways. Other pundits have announced that the business corporation, government institutions, and private organizations in general, along with the city and the nation-state are on their way to extinction. Familiar, cherished technologies such as the book, the newspaper, or the library appear to many to be scheduled for retirement, as are less attractive social features such as class, privilege, and economic disadvantage. Even quite abstract notions such as distance have had their obituaries read. Now, it seems, humanity may be adding itself to the list of the doomed.

One marked difference between the conventional digital endism and Joy's warning is that the former usually falls into the optimists' camp, while the latter is on the side of the pessimists. Here again is the old opposition of euphoria and gloom that we mentioned above. So, for example, when Victor Hugo's archdeacon of Notre Dame famously waved his book at the cathedral and cried,

> This will kill that. The book will kill architecture.

he reflected a fear that the public awe inspired by religious architecture would lose its force in face of the private aesthetic of the book, and he mourned the loss. When futurists type on their keyboard,

> This will kill that. The computer will kill the book.

the argument may seem quite similar (one technology will replace another), but these writers are now triumphant. The computer, they believe, should kill such things. The Joy debate also follows the same logic, but switches the evaluation back to the archdeacon's side. Genetic engineering, nanotechnology, and robotics are to be feared; what they will eliminate (principally humanity) is to be mourned.[14]

These periodic shifts between euphoria and gloom, then—however opposed their polarities—rely on a common logic that is independent of the human evaluation laid on top. This is the logic of technological determinism: the belief, whether held with hope or with resignation, that technology will determine social outcomes and so, by looking at the technology, you can predict what will happen and then decide whether you like it or not. This narrow sense of causality is the source of the tunnel vision we mentioned earlier, and it usually provides far too narrow a focus to make useful predictions.

Indeed, the first thing to note about such technological determinism is that it is often simply wrong. The book, the city, privilege, and the nation-state still thrive. For all its marvelous capabilities, digital technology is having difficulty pushing them out of the ring in what pundits predicted would be an easy and short contest. Take, as particular example, predictions about the office. For years, there has been talk about "electronic cottages." Technology, it is believed, will allow entrepreneurs to leave the drudgery of the high rise and set up on their own, changing cities, buildings, and economic organization radically. When the Internet came along, many assumed that, capable as it is of enhancing communication and reducing transaction costs, it would be the critical force to bring this transformation about in quick time. As a result, the issue provides a reasonable and testable case of current technology predictions at work.[15] The

Bureau of Labor Statistics regularly measures self-employment and so provides good evidence to judge progress to this vision. In 2000, the Bureau reported that 1994-1999, the first period of extensive Net connectivity (and the boom years before the dot com crash), was actually the first five-year span since 1960 in which the number of non-agricultural self-employed fell. People are not leaving organizations to set up on their own. They are leaving self-employment to join organizations. The reasons why are complicated, but the Bureau's study at least suggests that simple technological extrapolations, even in the most propitious of circumstances, can be as wrong as they are beguiling.[16]

Similarly, to return to where we began, the fatal narrowness of technological determinism helps explain why Langer's predictions of nuclear-determined social transformation did not come true "*in* [*his*] *lifetime*" or since. First, he failed to see other forces at work—forces that lay outside the tunnel described by his logic of technology. And consequently, second, he made simple extrapolations from the present to the future state of technologies, when almost every technologist knows that development and the steps from invention to innovation and use are usually anything but simple and predictable. Let us look at each of these problems for technological determinism in a little more detail.

Outside the tunnel

Even had the technological path to nuclear power been smooth, progress was never a straightforwardly technological matter with its path determined by a purely technological logic (whatever that may be). Geo-political concerns, military interests, scientific limitations, and fiscal constraints, for example, complicated decisions to adopt this technology from the start. Even envy played an important part. Indeed, the *Commissariat à l'Energie Atomique* in France (one of the countries most enthusiastic about nuclear power) admitted that many countries adopted nuclear power primarily because others had. Such keeping-up-with-the-neighbor envy was unlikely to lead to simple scientific, technological, or economic decision making. At the same time, decision-makers had to confront unforeseen technological problems, social concern, protest, and action—whether spurred by broad environmental worries, more narrow NIMBY fears, or the interests of gas and coal producers. Problems, concern, protest, and action in turn precipitated political intervention and regulation. Increasingly complex loops developed. None of these has a place in Langer's technologically determined vision of the future.[17]

The idea that social forces are at play in technology development, promoting some directions and inhibiting others, is hardly news. Indeed, the literature is so vast we shall barely address it here. On the one hand, there's a large body of socio-technical studies. On the other, there's the economic literature concerning welfare economics, externalities, and network effects and the way these (rather than some idea of inherent technological superiority) shape what technologies are developed, adopted, or rejected.[18] Nevertheless, some futurists continue to focus on technology in splendid isolation, as if technologies can be understood in their "own" terms and in isolation from the embedding social system.

So, without engaging this debate in full, we want to insist that technological and social systems are interdependent, each shaping the other. Gunpowder, the printing press, the railroad, the telegraph, and the Internet certainly shaped society quite profoundly. But equally, social systems, in the form of polities, governments, courts, formal and informal organizations, social movements, professional networks, local communities, market institutions, and so forth, shaped, moderated, and redirected the raw power of those technologies. The process resembles one of "co-evolution," with technology and society mutually shaping each other.[19] In considering one, then, it's important to keep the other in mind. Given the crisp edges of technology and the fuzzy ones of society, it certainly isn't easy to grasp the two simultaneously. Technological extrapolation can seem relatively easy. What Daniel Bell calls "social forecasting" is much harder.[20] But grasp both you must, if you want to see where we are all going or design the means to get there. And to grasp both, you have to reach outside the tunnel in which designs usually begin their life.

One small step...

What we tend to get, however, is the simpler kind of extrapolation, where the path to the future is mapped out along vectors read off from technology in isolation. Following these vectors, it's easy to count in the order of "1, 2, 3, … one million," as if intervening steps could be taken for granted. This way of counting occurs throughout the predictions of the nuclear age. A post-bomb book from 1945, written when no one had even developed a nuclear car engine, notes that

> Production of the atomic-energy type of motor car will not entail very difficult problems for the automobile manufacturer … it will be necessary to replace the 30,000,000 now registered in a few years.

Elsewhere, nuclear energy experts were predicting as late as 1974 that, spurred by the oil crisis, some 4,000 U.S. nuclear power plants would be on line by the end of the century. (The current figure around 100, with no new ones in production.) And Langer strides from the bomb to the U-235-powered house with similar ease, claiming blithely, "None of the things mentioned has yet been worked out, but the difficulties are difficulties of detail." Unbeknownst to him, the devil was clearly in these details, as he so often is.

Having shown some extraordinary bursts of exponential growth, digital technologies are understandably prey to this sort of extrapolation. Unfortunately, where real growth of this sort does occur (as with the explosion of the World Wide Web on the release of Mosaic, for example), it is rarely predicted, while where it is predicted, it often fails to occur. We are still waiting, as we have been for some decades, for the forever-just-around-the-next-corner natural language processing or strong artificial intelligence. It is always wise to recall that the age of digital technology has given us the notion of "vaporware." The term can embrace both product announcements that, despite all good will, fail to result in a product and announcements deliberately planned to prevent rivals from bringing a product to market. The forecasts of the digital age can, then, be a curious mixture of naïve predicting and calculated posturing. Both may be quite misleading. Yet they can also be informative. In understanding why intended consequences don't come about, we may find resources to fight unintended ones.

The road ahead

We have made this excursion into technological determinism, to throw some light of our own on Joy's three main areas of concern, bioengineering, nanotechnology, and robotics. On the one hand, determinism may easily count society out by counting in the fashion we described above. And on the other hand, determinism with its blindness to social forces excludes areas where the missing brakes on technology might be found before society autodestructs, a victim to its lust for knowledge. Let us contemplate the road ahead for each of these concerns, before looking at the theme, common to them all, of self-replication.

Bioengineering

By the late 1990s, "biotech" seemed an unstoppable force, transforming pharmaceuticals, agriculture, and ultimately mankind. The road ahead appeared to involve major chemical and agricultural interests barreling unstoppably along an open highway. Agricultural problems will be solved forever, cried the optimists. The environment will suffer irreparable harm, cried the pessimists. Many from both sides accepted that this future was inevitable. Within a remarkably short time, however, the whole picture changed dramatically. In Europe, groups confronted the bioengineering juggernaut with legal, political, and regulatory roadblocks. In India, protestors attacked companies trying to sell bioengineered seeds. In the United States, activists against genetically modified foods from around the world gathered to stop the WTO talks. Others attacked the GATT's provisions for patenting naturally occurring genes. Meanwhile, Monsanto has been forced to suspend research on sterile seeds. Carghill faces boycotts in Asia. Grace has faced challenges to its pesticides. And Archer Daniel Midlands has had to reject carloads of grain in the fear that they may contain traces of StarLink (a genetically engineered corn containing Cry9C protein, which may produce allergic reactions). Farmers, once committed to modified seeds have uprooted crops for fear that their produce will also be rejected. Around the world, champions of genetically modification, who once saw an unproblematic and lucrative future, are scurrying to counter consumer disdain for their products. If, as some people fear, genetic engineering represents one of the horses of the Apocalypse, it is certainly no longer unbridled nor its road ahead clear. The now-erratic biotech stocks remind those who bought them at their earlier highs how hard it is to extrapolate from current technology to the future.

As to that future, there's no clear consensus.[21] Euphoric supporters have plunged into gloom. "Food Biotech," one supporter recently told the *New York Times* gloomily, "is dead. The potential now is an infinitesimal fraction of what most observers had hoped it would be."[22] What does seem clear is that those who support genetic modification will have to look beyond the labs and the technology if they want to advance. They need to address society directly—not just by labeling modified foods, but by engaging public discussion about costs and benefits, risks and rewards. A little while ago, the same might have been said of the bioengineering as was said of nuclear technology: "decisions have been left to restricted groups, which were for a long time fairly insensitive to political trends within society."[23] Now, those trends are starting to become clear, while that insensitivity has extracted a heavy price. (The licensing of StarLink garnered Aventis CropScience less than $1 million; dealing with StarLink-contaminated crops is costing the firm hundreds of millions.) With interests other than the technologists' and manufacturers' involved, the nature of the decisions to be made has shifted dramatically from what can be done to what should be done. Furthermore, debates once focussed on biotechnologically determined threats now embrace larger issues concerning intellectual property in genetic code, distributive justice, preservation of biodiversity and a host of other socio-technological questions.[24] Of course, having ignored social concerns in a mixture described by one Monsanto scientist as "arrogance and incompetence ... an unbeatable combination," proponents have made the people they now must invite to these discussions profoundly suspicious and hostile.[25] In consequence, to make any progress on their envisioned road ahead, they first face a significant uphill drive.

Fears of bioengineering running rampant and unchecked are certainly fears to worry about. No one should be complacent. But if people fail to see the counteracting social and ethical forces at work, they end up with the hope or fear that nothing can be done to change the technologically determined future. But much is being done. Politicians, regulators, and consumers are responding to dissenting scientists and social activists and putting pressure on producers and innovators, who in turn must check their plans and activities. Looking at the technology alone predicts little of this, only the euphoria or the gloom. Undoubtedly, even with these constraints, mishaps, mistakes, miscalculations—and deliberate calculations—are still threats. But they are probably not threats that endanger humanity's existence. To keep the threat within bounds, it is important to keep its assessment in proportion.

Nanotechnology

If biotechnology predictions suffer from a certain social blindness, nanotechnology suffer from the alternative problem of thinking in the tunnel—too rapid extrapolation. Nanotechnology involves engineering at a molecular level to build artifacts "from the bottom up." Both the promise and the threat of such engineering seem unmeasurable. But they are unmeasurable for a good reason. The technology is still almost wholly on the drawing board. At Xerox PARC, Ralph Merkle, working with Eric Drexler, built powerful nano-CAD tools and then ran simulations of the resulting designs. The simulations showed definitively and in the face of skepticism that nano devices are theoretically feasible. This alone was a remarkable achievement. (And it emphasizes the importance of fundamental research, which by pushing at theoretical boundary conditions can open new regions for research and development.) But theoretically feasible and practically feasible are two quite different things.[26] It is essential not to leap from one to the other, as if the magnitude of the theoretical achievement made the practical issues ahead inconsequential (a mistake that, as we have seen, nuclear scientists like Langer made all too easily.) As yet, no one has laid out a detailed route from lab-based simulation or simple, chemically constructed nano devices to practical systems development.

So here the road ahead proves unpredictable not because of an unexpected curve, such as those genetically modified foods met, but because the road itself still lacks a blueprint. In the absence of a plan, it's certainly important to ask the right questions. Can nanotechnology actually fulfill its great potential in tasks ranging from data storage to pollution control? And can it do such things without itself getting out of control? But in fearing that nano devices will run amok, we are in danger of getting far ahead of ourselves. If the lesson of biotechnology means anything, however, even though useful nano systems are probably many years away, planners would do well to consult and educate the public early on.[27] And in fact, the proponents of nanotechnology are doing that. Eric Drexler raised both benefits and dangers in his early book *Engines of Creation* and has since founded the Foresight Institute to help address the latter. Following the National Institute of Health's (NIH) example in the area of recombinant DNA research, the Foresight Institute has also created a set of guidelines for research into nanotechnology.[28]

Robotics

Robots, popularized by science-fiction writers such as Arthur C. Clarke and Isaac Asimov and now even commercialized as household pets, are much more familiar than nanodevices. Nonetheless, as with nanotechnology, the road ahead, whether cheered or feared, has appeared in our mapbooks long before it will appear on the ground. Again many of the promises or problems foreseen show all the marks of tunnel vision.[29] Critical social factors have been taken for granted; profound technological problems have been taken as solved. The steps from here to a brave new world have been confidently predicted, as if all that remained was to put one foot in front of the other. 2001, after all, was to be the year of Hal. Hal will have to wait a while yet to step from the screen or page into life.

Take for example the cerebral side of the matter, the much-talked about "autonomous agents" or "bots." These are the software equivalent of robots, which search, communicate, and negotiate on our behalf across the Internet. Without the impediment of a physical body (which presents a major challenge for robotics), bots, it has been claimed, do many human tasks much better than humans and so represent a type of intelligent life that might come to replace us. Yet bots are primarily useful because they are quite different from humans. They are good (and useful) for those tasks that humans do badly, in particular gathering, sorting, selecting, and manipulating data. They are contrastingly often quite inept at tasks that humans do well—tasks that call for judgement, taste, discretion, initiative, or tacit understanding. Bots are probably better thought of as complementary systems, not rivals to humanity. Consequently, though they will undoubtedly get better at what they do, such development will not necessarily make bots more human or rivals for our place on earth. They are in effect being driven down a different road. Certainly, the possibility of a collision between the decision-making of bots and of humanity needs to be kept in mind. In particular, we need to know who will be responsible when autonomous bots inadvertently cause collisions—as well they might. The legal statuses of autonomous actors and dependent agents are distinct, so autonomous agents threaten to blur some important legal boundaries. But we probably need not look for significant collisions around the next few bends. Nor, should they come, should we expect them to threaten the very existence of humanity.

Are more conventional, embodied robots—the villains of science fiction—any greater threat to society? We doubt it, even though PARC research on self-aware, reconfigurable polybots has pushed at new robotic frontiers. These, combined with mems (microelectical mechanical systems), point the way to morphing robots whose ability to move and change shape will make them important for such things as search and rescue in conditions where humans cannot or dare not go. Yet, for all their cutting-edge agility, these polybots are a long way from making good free-form dancing partners. Like all robots (but unlike good dancing partners), they lack social skills. In particular, their conversational skills are profoundly limited. The chatty manner of C3-PO still lies well beyond machines. What talking robots or computers do, though it may appear similar, is quite different from human talk. Indeed, talking machines travel routes designed specifically to avoid the full complexities of situated human language. Moreover, their inability to learn in any significant way hampers the proclaimed intelligence of robots. (This failing has apparently led Toyota, after heavy investment in robotics, to consider replacing robots with humans on many production lines.) Without learning, simple common sense will lie beyond robots for a long time to come. Indeed, despite years of startling advances and innumerable successes like the chess-playing Big Blue, computer science is still almost as far as it ever was from building a machine with the learning abilities, linguistic competence, common sense, or social skills of a five year old.

So, like bots, robots will no doubt become increasingly useful. But given the tunnel design that often accompanies tunnel vision, they will probably remain frustrating to use and so seem anti-social.[30] But (though the word *robot* comes from a play in which robots rebelled against their human masters) this is not anti-social in the way of science fiction fantasies, with robot species vying for supremacy and Dalek armies exterminating human society. Indeed, robots are handicapped most of all by their lack of a social existence. For it is our social existence as humans that shapes how we speak, learn, think, and develop common sense and judgement. All forms of artificial life (whether bugs or bots) are likely to remain primarily a metaphor for—rather than a threat to—society at least until they manage among themselves to enter a debate, form a choir, take a class, survive a committee meeting, join a union, design a lab, pass a law, engineer a cartel, reach an agreement, or summon a constitutional convention. Such critical social mechanisms allow society to shape its future, to forestall expected consequences (such as Y2K), or to respond to unexpected ones (such as epidemics).

Self-replication

As we noted earlier, one critical, pervasive concern runs through each of these examples: the threat from self-replication. The possibility of self-replication is most evident with bioengineering, as the biological organisms on which it is based are already capable of reproducing (although bioengineering has deliberately produced some sterile seeds). Both nanotechnologists and robotic engineers are also in pursuit of self-replicating artificial life.[31] So it is certainly reasonable to fear, that, in their ability to replicate, bioengineered organisms or mechanical devices may either willfully (as intelligent robots) or blindly (as nano devices or organisms) overrun us all.

Despite the apparently unprecedented nature of these sciences, the threat of self-replication to humanity itself has a history. The problem was first laid out in detail by the eighteenth-century economist, Thomas Malthus (who drew on previous work by, among others, Benjamin Franklin). Two hundred years ago in his famous *Essay on the Principle of Population*, Malthus argued that the self-replicating threat to humanity was humanity itself, and the nature of the threat was susceptible to mathematical proof. Population, he claimed, grew geometrically. Food production only increased arithmetically. The population, therefore, would inexorably outstrip its means of support. By 2000, he extrapolated, the population would be 256 times larger, while food production would only have grown nine-fold. His solution was to slow the growth of the population by making the environment in which the poor reproduced so unbearable that they would stop reproducing. Attempts to ameliorate poverty, in his view, threatened humanity itself.

In defining the problem and hence designing a solution, Malthus—and the poor—were victims of his tunnel vision. He extrapolated from present to future as if the iron laws of mathematics bound the growth of society and its resources. In fact, in the nineteenth century, agricultural production and productivity increased dramatically, while the shift, with industrialization, to urban life reduced both the need for and social norm of large families. Worldwide, the population has grown only grown six fold not 256 fold since Malthus's time.[32] Growth in productivity and in land under production has kept pace. No one should underestimate the critical issues of diminishing returns to agricultural production (through pollution and degradation) or of equitable distribution, but these issues fall mostly outside Malthus's iron law and iron-hearted solution. Malthus, then, was a victim once again both of tunnel vision (he saw only a restricted and predictable group of forces at work) and of over eager counting. His argument also shows how misunderstanding the scale of a threat to humanity can lead to inhumane responses. "The Malthusian perspective of food-to-population ratio," the Nobel-laureate economist Amartya Sen notes, "has much blood on its hands."[33]

Now the threat of replication appears from a different quarter. And rather than humans outstripping technology, it seems to be technology that is outstripping humanity. Nonetheless, the problem may look quite as inescapable as it did to Malthus. Bioengineered organisms, nano devices, and robots, might quite literally take on and sustain a life of their own, leading with Malthusian inexorability to a future that "doesn't need us." Molecular biology might produce a "white plague." Replicating nano devices might reproduce unstoppably, damaging the molecular structure of our world imperceptibly. The damage may only be evident when it is too late to respond. Without intelligence or intention, then, either may blindly eliminate us. On the other hand, "once an intelligent robot exists," Joy fears, "it is only a small step to a robot species"—a species that may first outsmart us and then quite deliberately eliminate us.[34] Let us take the robots first. Clearly, we have doubts about such claims for intelligence. Leaving those aside, we have even graver doubts about that "small step." Here we are not alone. At the "Humanoids 2000" conference at MIT, experts in the field were asked to rank on a scale of zero to five the possibility that robots "will eventually displace human beings." Their collective wisdom rated the possibility at zero. For most, the "small step" would need something well beyond the seven-league boots of legend.[35]

But what of unintelligent replication? Undoubtedly, as examples from kudzu to seaweed remind us, when replicating organisms find a sympathetic niche with no predators, they can get out of hand very quickly. From aquarium outlet pipes, Australasian taxifolia is spreading over the sea floor off the coast of France and California. It's not threatening human life itself, and unlikely to, but it offers a clear example of rampant self-replication. (Some scientist suspect that the Mediterranean variety is a mutant, whose genes were changed by the ultraviolet light used in aquariums.) Yet even here, it's important to note, self-replication is not completely autonomous. It depends heavily on environmental conditions—in the case of the seaweed, the sympathetic niche and absence of predators, while the species under threat, principally other seaweeds, are incapable of collective, corrective action. These considerations are even more important in considering the threat of self-replicating artificial life, which is highly sensitive to and dependent on its environment. "Artificial self replicating systems," Ralph Merkle notes "function in carefully controlled artificial environments." They are simply not robust enough to set up on their own and beyond our control.[36]

New organisms and devices, then, do not exist and will not replicate in a world of their own making. Replication is an interdependent process. This fact doesn't minimize the gravity of the accidental or malicious release from a lab of some particularly damaging organism. But it does suggest how we might minimize the danger, reminding us of physical and social resources at our disposal—resources that technologically narrow tunnel vision makes it hard to see. For while tunnel vision views technologies in isolation, ecology insists that complex, systemic interdependence is almost always necessary for reproduction. Larger ecosystems as a whole need to reproduce themselves in order for their dependent parts to survive, and vice versa. Within such environments, sustainable, chain reactions live in a fairly narrow window. If they are too weak, they are a threat to themselves; if they are too strong, they threaten the ecosystem that supports them. This window is evident in nuclear reactions. If they are too far below critical, the reaction is not sustained. If they are too far above it, the fuel is destroyed in a single unsustainable explosion. Similarly, organic viruses have to be efficient enough to survive, yet the first lesson for any new virus is simply "Don't be too efficient": if it is, it will kill its host and so destroy its reproductive environment. When humanity is the host, viruses face an extra problem. In this case, unlike in the case of native Mediterranean seaweeds, when the host survives, it can organize collectively to combat the virus that is killing it individually.

So we have to look at the threat of replication in terms of environmental factors and the social and institutional organization, rather than in terms of the organism, nano device, or robot on its own. History suggests that, directly or indirectly, humanity has been good at manipulating environments to limit replication. Responses to destructive organisms have included informal and formal institutional standard-setting. These stretch from social norms (such as those that fostered hand washing or eradicated spitting in public—Washington, after all, was described by Dickens as the "capitol of expectoration," but even it has managed to improve) to sanitary and health codes (dealing with everything from burial places to reporting and tracking communicable diseases). Similarly, through the institutions and norms of network etiquette, ranging from categorically forbidden practices to widely accepted standards, people deal with computer viruses. Certainly, computers are always under threat from viruses, while information warfare threatens whole networks. But the user's control over the environment (which ultimately includes the ability to disconnect from networks or simply shut down, but has many other resources for protection before reaching that point) provides a powerful, environmental countermeasure that make attempts to annihilate computing through blind replication less plausible than they might seem.[37]

We do not want to oversimplify this claim for environmental protection of humanity (or computer networks). First, we must stress we are not embracing a naive "Gaia" hypothesis and claiming that the environment is sufficiently self-adaptive to handle all threats. It is not. Rather, we are claiming (and we expand on this point below) that in order to deal with self-replicating threats, society can organize to adjust the environment in its own defense. But second, we must also stress that this is not an unproblematic solution. Malthus, after all, wanted to adjust the environment in which the poor lived. And playing with the environment is fraught with difficulties. It can be hard to cleanse a biological environment without destroying it (a profound problem in treating cancer patients). And it can be hard to see the collateral damage such strategies may give rise to. DDT, for example, while effectively destroying the environment of malaria-carrying insects, did long-term damage to other species. Indeed, intermittently society has proved lethally good at destroying sustainable environments, producing inert deserts and dust bowls in their place. But, luckily, it is probably less proficient at building robust, sustainable environments—particularly ones over which it will simultaneously lose control, enabling the organisms within the environment to wipe humanity out. Yet such an environment is a precondition for an all-out, society-destroying threat from self-replicating nanotechnology or robots. These will emerge weak and vulnerable on the path to self-replication, and thus they will be heavily dependent on, and subject to control through the specialized environments that sustain them. “It is difficult enough,” Merkle acknowledges, “to design a system able to self replicate in a controlled environment, let alone designing one that can approach the marvelous adaptability that hundreds of millions of years of evolution have given to living systems.”[38]

Again it is important to stress that threats are real. Society may well cause a good deal of destruction with these technologies as it has with nuclear technology. And it may also do a lot of damage attempting to defend itself from them—particularly if it overestimates the threat. But at present and probably for a good while into the future, the steps from current threat to the annihilation of society that Joy envisages are almost certainly harder to take than the steps to the control and contain such a threat. Blindly or maliciously, people may do savage things to one another and we have no intention of making light of the extent of such damage if amplified by current technologies. But from the current threat to the destruction of society as a whole (rather than just a part of it) may be less like the gap from one to one million and more like the gap between one million and infinity—of a different order of magnitude entirely. Furthermore, as the example of Malthus reminds us, exaggerated threats can lead to exaggerated responses that may be harmful in themselves. Malthus’s extrapolations provided a rational for the repressive Poor Laws of the nineteenth century. Overestimating the threats that society faces today may in a related fashion provide a climate in which civil liberties are perceived as an expendable luxury. Repressive societies repeatedly remind us that overestimating threats to society can be as damaging as underestimating them.

Self organization

An essential feature—perhaps a uniquely human one—in these responses to threats from self-replication is organization. In anticipation of or in response to threats, humans organize themselves in a variety of ways. Determined attempts to change environments in which malign bacteria can replicate have usually demanded determined organization of one form or another. Today, to preempt threats from bioengineered organisms, for instance, governments and bodies like the NIH help to monitor and regulate many labs whose work might pose an overt risk to society. Independent organizations like the Foresight Institute and the Institute for Molecular Manufacturing, as we have seen, also attempt to ensure responsible work and limit the irresponsible as much as possible.[39] Such guidelines rely in part on the justifiable assumption that labs are usually part of self-organizing and self-correcting social systems. For labs today are not the isolated cells of Dr. Frankenstein. Rather, they are profoundly interrelated. In biotechnology, some of these relations are highly competitive, and this competitiveness itself acts to prevent leaks, deliberate or accidental. Others are highly interlinked, requiring extensive coordination among labs and their members. (Some papers in the field have more than 100 authors.) In these extended networks, people are looking over each other’s shoulders all the time. Not only does this constant monitoring provide important social and structural limits on the possibility of releases and the means to trace such releases to their source. It also distributes the knowledge needed for countermeasures.[40]

There are many similarities in the way that the Net works to ward of software viruses. Its interconnectivity allows computer viruses to spread quickly and effectively. But that very interconnectedness also helps countermeasures spread as well. Supporting both self-organizing and intentionally organized social systems, the Net allows the afflicted to find cures and investigators to track sources. It also creates transient and enduring networks of people who come together to fight existing threats or stay together to anticipate new ones. In the labs and on the Net, as we can see, the systems that present a threat may simultaneously create the resources to fight it.[41] But society cannot use these networks creatively if tunnel vision prevents it from seeing them. To repeat, we do not want to diminish the seriousness of potential attacks, whether in wetware or software, whether intentional or accidental. We do, however, want to bring the threat into proportion and prevent responses from getting out of proportion (whether in the direction of Malthusian savagery or Luddite despair). Replication is not an autonomous activity. Thus control over the environment in which replication takes place provides a powerful tool to respond to the threat.

Demystifying

The path to the future can look simple (and sometimes simply terrifying) if you look at it through tunnel vision—or what we have also called 6-D lenses. We coined this phrase having so often come upon “de-” or “di-” words like *demassification*, *decentralization*, *disintermediation*, *despacialization*, *disaggregation*, and *demarketization* in futurology. These are grand technology-driven forces which some futurists see spreading through society and unraveling our social systems. If you take any one of the Ds in isolation, it’s easy to follow its relentless journey to a logical conclusion in one of the endisms we mentioned earlier. So, for example, because firms are getting smaller, it’s easy to assume that firms and other intermediaries are simply disintegrating into markets of entrepreneurs. And because communication is growing cheaper and more powerful, it’s easy to believe in the death of distance. But these Ds rarely work in such linear fashion. Other forces (indeed, even other Ds) are at work in other directions. Some, for example, are driving firms into larger and larger mergers to take advantage of social (rather than just technological) networks. Yet others are keeping people together despite the availability of great communications technology. So, for example, whether communications technology has killed distance or not, people curiously just can’t stay away from the social hotbed of modern communications technology, Silicon Valley.[42]

To avoid the mistake of reading the future in such a linear fashion, we need to look beyond individual technologies and individuals in isolation. Both are part of complex social relations. And both offer the possibility of social responses to perceived threats. Thus looking beyond individuals offers alternatives to Joy’s suggestion that the best response to potentially dangerous technologies is for principled individuals to refuse to work with them. Indeed, Joy’s own instinctive response to the threat he perceived, as his article makes clear, was to tie himself into different networks of scientists, philosophers, and so on.[43] This point is worth emphasizing because to a significant degree the digerati have been profoundly individualistic, resisting almost any form of institution and deprecating formal organizations while glorifying the new individual of the digital frontier. This radical individualism can, as we have been suggesting, lead both to mischaracterizing the problems society faces and overlooking the solutions it has available to respond. In particular, it tends to dismiss with contempt any forms of institutional response. Institutions are, it can seem, for industrial age problems and worse than useless in the digital age.

Undoubtedly, as the Ds indicate, old ties that bound communities, organizations, and institutions are being picked at by technologies. A simple, linear reading then suggests that these will soon simply fall apart and so have no role in the future. A more complex reading, taking into account the multiple forces at work, offers a different picture. Undoubtedly particular communities, organizations, and institutions will disappear. But communities, organizations, and institutions *sui generis* will not. Some will reconfigure themselves. So, while many nationally powerful corporations have shriveled to insignificance, some have transformed themselves into far more powerful transnational firms. And while some forms of community are dying, others bolstered by technology are being born. The virtual community, while undoubtedly overhyped and mythologized, is an important new phenomenon.[44]

Undoubtedly, too, communities, organizations, and institutions can be a drag on change and innovation. But for this very reason, they can act to brake the destructive power of technology. Delay, caution, hesitation, and deferral are not necessarily bad, particularly when speed is part of the problem.[45] Moreover, communities, organizations, and institutions have also been the means that has given us technology's advantages. Scientific societies, universities, government agencies, laboratories, not lone individuals, developed modern science. As it continues to develop, old institutional forms (copyright and patent law, government agencies, business practices, social mores, and so forth) inevitably come under stress. But the failure of old types of institution is a mandate to create new types.

Looking back, Robert Putnam's book *Bowling Alone* shows the importance of institutional innovation in response to technological change. The late nineteenth century brought the United States profound advances in scientific research, industrial organization, manufacturing techniques, political power, imperial control, capital formation, and so on. These accompanied unprecedented technological advances, including the introduction of cars, airplanes, telephones, radio, and domestic and industrial power. They also brought migration and social deracination on an unprecedented scale. People moved from the country and from other countries to live in ill-prepared, ill-built, polyglot, and politically corrupt cities. Social disaffection spread as rapidly as any self-replicating virus—and viruses spread widely in these new, unsanitary urban conditions too. The very social capital on which this advanced society had built was devalued more quickly than any Internet stock. There was, Putnam notes, no period of economic stress quite like the closing years of the old century.[46]

But in response, Putnam notes, the early years of the new century became a remarkable period of legal, government, business, and societal innovation—stretching from the introduction of anti-trust legislation to the creation of the American Bar Association, the ACLU, the American Federation of Labor, the American Red Cross, the Audubon Society, 4H, the League of Women Voters, the NAACP, the PTA, the Sierra Club, the Urban League, the YWCA, and many other associations. Society, implicitly and explicitly, took stock of itself and its technologies and acted accordingly. The resulting social innovation has left marks quite as deep as those left by technological innovation, rebuilding social capital, rekindling innovation, removing old institutional road blocks and replacing them with new institutions.

The dawn of the atomic age offers another precedent. For all its difficulties in the laboratory proper, the economist Michael Damian suggests that nuclear power nonetheless did create a powerful and instructive laboratory, a social laboratory that confronted with an extraordinary range of problems, interrelated as never before,

> In terms of risk, societal issues, and the democratic accountability of industrial choice; in terms of control over externalities, with the interlocking of political, social and biological dimensions in economic issues; in terms of megascience, complexity and technology in extreme or hostile environments; more generally, in terms of the relations between work and life, and the relations between nations.[47]

Out of this confrontation came the new national and international institutions of the atomic age to address major problems from nuclear proliferation to nuclear waste disposal. These included the Atomic Energy Commission, the Committee for Nuclear Disarmament, and the Nuclear Regulatory Commission. Faced once more with unprecedented change, we need similar social laboratories to those of the 1950s, similar inventiveness to the 1900s. We are not calling here for the perpetuation of old institutions. Indeed, Putnam’s work suggests that many of these are fading rapidly. Rather, we see in Joy’s nightmare a need for the invention of radically new institutions. Energy equivalent to that which has been applied to technology needs to be applied to the development of fitting institutions. Moreover, we see the need for people to realize that, like it or not, new institutions are coming into being (whatever the hype about an institutionless future), and it is increasingly important to develop them appropriately to meet the real challenges that society faces.[48]

In claiming that, outside technology’s tunnel, we can see brakes that Joy can’t, we are not, we should emphasize once again, encouraging complacency. To prevent technologies getting destructively out of control will require careful analysis, difficult collective decision making, and very hard work. Our goal is to indicate where the work is needed. As in the progressive era, as with the nuclear case, so we believe in the era of digital technologies, social resources can and must intervene to disrupt the apparently simple and unstoppable unfolding of technology and to thwart the purblind euphoria and gloom of prognosticators. Intriguingly, this sort of social action often makes judicious commentators like Oppenheimer become curiously self-*un*fulfilling prophets. Because they speak, because they raise awareness, society is able to prevent the doom they prophesy from coming about. As we have tried to argue, society cannot ignore such prophecies because previous examples have turned out to be wrong. Rather, it must take the necessary action to prove them wrong. If people do not freeze in terror at his vision, Bill Joy will, we hope and believe, be proved wrong. Society will respond to the threat and head off disaster. But, paradoxically once again, this will only make raising awareness, all the more right.

---

[1] R.M. Langer: "Fast New World" *Colliers National Weekly*, July 6, 1940: 18-19 & 54-55.

[2] Here at least Dr. Langer was not out on a limb. Experts in their fields made similar predictions. William B. Stone, past president of the Society for Automobile Engineers, was confident that "motor vehicles will be propelled with this [nuclear] energy." The National Advisory Committee for Aeronautics predicted planes with atomic engines. Ships, it was said, would cross the Atlantic on a spoonful of fuel. See *The Atomic Age Opens*. New York: Pocket Books, 1945.

[3] John J. O'Neill, *The Almighty Atom: The Story of Atomic Energy* (New York: Washburn, Inc., 1945). S.H. Schurr and J. Marschak, *Economic Aspects of Atomic Power* (Princeton, NJ: Princeton University Press for Cowles Commission for Research in Economics and Science, 1950); President Dwight D. Eisenhower, "Atoms for Peace" speech presented to the United Nations, December 8, 1953. On the egalitarian prospects of nuclear power, it is worth noting that today, 90 percent of nuclear power is generated in a dozen countries; except for Russia, all of these are OECD members.

[4] See Dave Walter, *Today Then: Americans Best Minds Look 100 Years into the Future on the Occasion of the 1893 World's Columbian Exposition* (Helena, MT: American and World Geographic Publications, 1992).

[5] Utopian visions often accompany warfare and terror, either encouraging protagonists in advance or easing their consciences retrospectively. See R. Schaer, G. Gaeys, and L. Sargent (eds.), *Utopia: The Search for the Ideal Society in the Western World* (New York, NY: Oxford University Press for New York Public Library, 2000), in particular, Bronislaw Baczko, "Fin-de-Siècle Landscapes Against a Background of Ruins", ibid., pp. 201-205.

[6] J. Robert Oppenheimer, "The Atomic Age," in Hans Albrecht Bethe, Harold Clayton Urey, James Franck, J. Robert Oppenheimer, *Serving Through Science: The Atomic Age* (New York, NY: United States Rubber Company, 1945), p. 14.

[7] Bill Joy, "Why the Future Doesn't Need Us," *Wired* 2000 8.04: 238-262.

[8] It's surely this narrowness that allows Joy, a sensitive and ethical figure, to find his ideas inextricable from those of Ted Kaczynski, the nihilistic "Unabomber."

[9] Digerati hype, however constructively meant, may well be disarming in its effects. It seems reasonable to argue, for example, that euphoric account of the e-economy killed some robust business plans that might have survived the current downturn while cheering into their place unsustainable plans that have since collapsed catastrophically.

[10] Indeed, it is fairer to say that Joy brought an existing debate among specialists to a broader public with his article than that he initiated a debate. For a less trenchant view of the issues by someone who has been discussing them for a while see Neil Jacobstein, "Values-Based Technology Leadership and Molecular Nanotechnology," paper presented at the 50th Anniversary of the Aspen Institute, Aspen, CO August, 2000.

[11] That danger is more than imaginary. A single nuclear plant in Scotland recently failed to account for nearly 375 pounds of Uranium; 250 pounds more is missing in Lithuania.

[12] At Los Alamos, Edward Teller had the task of allaying these fears, which he tried to do by showing that their probability was extremely low. Despite his calculations, however, fears appear to have persisted right up to the first actual detonation. War Department report, quoted in *The Almighty Atom*, p. 51.

[13] For our discussion of endism, see John Seely Brown and Paul Duguid, *The Social Life of Information* (Boston, MA: Harvard Business School Press, 2000), especially chapter 1.

[14] Pessimism is not inevitable. There are optimists who appear to relish an envisaged end for humanity. Ray Kurtzweil seems to think we shall give up human bodies for silicon ones. John Perry Barlow apparently hopes that we may give up our bodies all together and upload our minds into cyberspace. Similarly, Kevin Kelley embraces the "Out of Control" state that worries Joy so profoundly.

[15] When asked "What kind of evidence would persuade you of fundamental change in the economy," the economist Paul Krugman, a skeptic about the Net's transformational powers, answered, "If I saw some really radical change, for example, in the typical form of economic organization - if we really did become a society of self-employed consultants forming temporary alliances doing business over the Internet." Kevin Kelly, "New Economy? What New Economy," *Wired* 1998 6.05. Online: http://www.wired.com/wired/archive/6.05/krugman.html.

[16] David Leonhardt, "Entrepreneurs' 'Golden Age' Has Failed to Thrive in 90s," *New York Time*, December 1, 2000, p. 1.

[17] Michael Damian, "Nuclear Power: The Ambiguous Lessons of History" *Energy Policy*, 20(7): 596-607, p. 598. See Stephen R. Barley, "Technology, Power, and the Social Organization of Work: Towards a Pragmatic Theory of Skilling and Deskilling," *Research in the Sociology of Organizations* 1988 1: 33-80 for other examples of keeping-up-with-the-neighbors as a motive for technological innovation.

[18] As an early example from the former, see W. Bijker, T. Hughes, & T. Pinch (eds.), *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* (Cambridge, MA: MIT Press, 1987). As a classic example of the latter, see Paul David, "Understanding the Economics of QWERTY: The Necessity of History" in William Parker (ed.), *Economic History and the Modern Economist* (Oxford, UK: Basil Blackwell). Since these early interventions, the literature in both fields has grown enormously.

[19] The notion of co-evolution is a tricky one. Douglas Engelbart, one of the great pioneers of modern computers, uses it to propose a series of (four) stages in the evolution of humanity. (See Thierry Bardini, *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing* (Palo Alto, CA: Stanford University Press, 2000), especially pp. 53-56.) These are powerful insights into the relationship between society and technology. They are not without problems, however. For example, Engelbart has also suggested that the two components of this evolution have recently pulled apart, so that "technology is erupting," while society is falling behind. Such a claim assumes that technology has an inherent developmental logic independent of society and can take on an evolutionary life of its own. Moreover, some versions of this argument echo Malthus's claims about population and succumb to the same objections (see below). As we argue later in this paper, evolutionary perspectives suggest that the components of a co-evolution are co-dependent. The idea that one part can suddenly become independent, autonomous, and self-replicating entails a profound evolutionary leap. (For further discussion of this point, see *The Social Life of Information*, pp. 84-5, where Engelbart's full quotation and its source are given.) What do tend to come apart are new technologies and old institutions framed around prior technologies and their uses. This leads to our conclusion that it is not humanity that needs reforming at times of technological change, but its institutions.

[20] Daniel Bell, *The Coming of Post-Industrial Society: A Venture in Social Forecasting* (New York, NY: Basic Books, 1973)

[21] A year-end 2000 report in the *New York Times* science section sums up this lack of clarity by opening: "More than a decade after the federal government allowed the first release of a genetically engineered organism into the environment, researchers concluded that scientists still cannot say with any precision what the ecological effects—either good or bad-of such genetically modified organisms might be." Carol Kaesuk Yoon, "What's Next for Biotech Crops? Questions. *New York Times*, December 19, 2000 C 1. The article reflects a longer report in *Science*.

[22] Dr. Henry Miller, senior research fellow at the Hoover Institute quoted in "Biotechnology Food: From the Lab to a Debacle," *New York Times*, January 25, 2001, p. C1.

[23] Damian, p. 601.

[24] For such debates, see, for example, James Boyle, *Shamans, Software, and Spleens: Law and the Construction of the Information Society* (Cambridge, MA: Harvard University Press, 1996) or Allen Buchanan, Dan Brock, Norman Daniels, and Daniel Wickler, *From Chance to Choice: Genetics and Justice* (New York, NY: Cambridge University Press, 2000). It is not simply that an age of innocence has given way to one of sophistication. Boyle argues that there was little innocence in the earlier age. Behind earlier, superficially technological claims, he sees a fundamental attempt to appropriate part of the information commons, an appropriation which promises rich rewards for the claim jumpers, but detrimental results for society. See also, Larry Lessig, *Code and Other Laws of Cyberspace* (New York, NY: Basic Books, 1999).

[25] Will Carpenter, former head of Monsanto's biotechnology strategy group, quoted in Biotechnology Food: From the Lab to a Debacle."

[26] See David Harel, *Computers, Ltd.: What They Really Can't Do* (New York: Oxford University Press, 2000) for the distinction between computable and intractable problems.

[27] The "public" here has to be broadly construed. In the case of bioengineering, the strongest resistance came from Europe, where the proponents PR campaigns and lobbying were weakest. For more on this, see *The Social Life of Information*, p. 251.

[28] Eric Drexler, *Engines of Creation* (Garden City, NY: Doubleday, 1986). For the NIH guidelines, see http://www.ehs.psu.edu/biosafety/nih/nih95-1.htm. For the Foresight Institute guidelines, see http://www.foresight.org/guidelines.

[29] Tunnel vision may also limit Ray Kurzweil's prognostications. As his article shows, these have heavily influenced Joy.

[30] For this notion of "tunnel design," see *The Social Life of Information*, pp. 2-4.

[31] For discussion of nanotechnological issues, see Ralph Merkle, "Self-Replication and Nanotechnology" online: http://www.zyvex.com/nanotech/selfRep.html. And for discussion of the robotic issues see Rodney Brooks, "Artificial Life: From Robot Dreams to Reality" *Nature* 2000 406: 945-947 and Hod Lipson and Jordan Pollack, "Automatic Design and Manufacture of Robotic Lifeforms" *Nature* 2000 406: 947-948.

[32] In the U.S., population has grown 55 fold since 1800, as the latest (2000) census data reveals, but here issues of territorial expansion and immigration complicate questions of self-replication, which is why the global figure is more informative.

[33] Malthus insisted that "There is no reason whatever to suppose that anything beside the difficulty of procuring in adequate plenty the necessaries of life should either indispose this greater number of persons to marry early or disable them from rearing in health the largest families." Consequently, he concluded, "The increase of the human species can only be kept down to the level of the means of subsistence by the constant operation of the strong law of necessity, acting as a check upon the greater power." His narrow argument would suggest that birth rates and family sizes would be lowest in poor countries and highest in rich ones. Roughly the opposite obtains today. Sen notes how Malthusian economists tend to focus on individual decision making and to ignore the "*social* theories of fertility decline." More generally, Sen argues that economists narrow their sources of information and ignore the social context in which that information is embedded. Amartya Sen, *Development as Freedom* (New York: Alfred A. Knopf, 1999).

[34] Part of Joy's reasoning here seems to be based on an article about replicating peptides. Even the author of this study is unsure whether what he has discovered is (in his own intriguing terms) "a mere curiosity or seminal," and it is certainly a very long step from peptides to a robot society. Stuart Kauffman, "Self-Replication: Even Peptides Do It," *Nature* 1996 382: 496-497.

[35] Kenneth Chang, "Can Robots Rule the World? Not Yet." *New York Times*, September 12, 2000, F1.

[36] "Spreading Tropical Seaweed Crowds Out Underwater Life" *St Louis Post-Dispatch*, October 26, 1997; Merkle. "Self Replication."

[37] We don't want to underestimate the threat to computer networks, on which economically advanced societies have become remarkably dependent, despite the networks' evident fragility. But threats to computer networks are not threats to the existence of humanity, and it is the latter which is the concern of "Does the Future Need Us?" and similar chiliastic visions.

[38] Merkle goes on,

> Designing a system that uses a single source of energy is both much easier to do and produces a much more efficient system: the horse pays for its ability to eat potatoes when grass isn't available by being less efficient at both. For artificial systems where we wish to decrease design complexity and increase efficiency, we'll design the system so that it can handle one source of energy, and handle that one source very well.

(Self Replication and Nanotechnology. Online: http://www.zyvex.com/nanotech/selfRep.html.) Such a system, like efficient viruses, will be equally fragile, depending on a single source for survival. Cut off its source of energy, and you cut of its ability to replicate.

[39] Apart from specific guidelines on positive behavior, the guidelines try to organize the network of scientists to isolate negative behavior. For example, the guidelines' principles include: "Governments, companies, and individuals who refuse or fail to follow responsible principles and guidelines for development and dissemination of MNT should, if possible, be placed at a competitive disadvantage with respect to access to MNT intellectual property, technology, and markets."

[40] For the social structure of biotechnology labs, see Karin Knorr Cetina, *Epistemic Cultures: How the Sciences Make Knowledge* (Cambridge, MA: Harvard University Press, 1999). For collaboration between labs, see Walter W. Powell, Kenneth W. Koput, and Laurel Smith Doerr, "Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology," *Administrative Science Quarterly*, 1996 41: 116-145.

[41] On the Net, it is worth noting, recent viruses have rarely been, as Robert Morris's infamous worm was, technologically driven. The famous "I love you bug," for example, was more socially innovative. It spread not by initiating its program autonomously, but by persuading people to run it. It did that by playing on human vanity and curiosity, which persuaded people to rush to open messages from distant acquaintances headed "I love you." But the social system of the Net also helps track down such hackers. Similarly, a central strategy in restricting the spread of the AIDS virus has been to change the practices that spread it, rather than to attack the virus itself.

[42] We elaborate this argument in John Seely Brown and Paul Duguid, "Mysteries of the Region: Knowledge Dynamics in Silicon Valley," pp.16-39 in C. Lee, W. Miller, M. Hancock, & H. Rowen (eds.), *The Silicon Valley Edge: A Habitat for Innovation and Entrepreneurship* (Stanford, CA: Stanford University Press, 2000).

[43] By contrast, as Joy's example of Ted Kaczynski suggests, individuals who cut themselves off from social networks can pose a significant threat (though hardly to society as a whole).

[44] See, for example, Barry Wellman and Milena Gulia, "Net Surfers Don't Ride Alone: Virtual Community as Community," pp. 331-367 in B. Wellman (ed.), *Networks in the Global Village* (Boulder, CO: Westview Press, 1999).

[45] For an elaboration of this argument about deferral and the digital age, see Carla Hesse, "Humanities and the Library in the Digital Age," pp.107-121 in Alvin Kernan (ed.), *What's Happened to the Humanities?* (Princeton, NJ: Princeton University Press, 1997)

[46] See Robert Putnam, *Bowling Alone: The Collapse and Revival of American Community* (New York: Simon and Schuster, 2000), especially chapter 23.

[47] Damian, 605.

[48] This is the kernel of Boyle's and Lessig's arguments (see above, fn. 24). See also the work of Phil Agre, in particular "Institutional Circuitry: Thinking about the Forms and Uses of Information" *Information Technology and Libraries* 1995 14 (4): 225-230; Walter Powell and Paul DiMaggio (eds.), *The New Institutionalism in Organizational Analysis* (Chicago, IL: Chicago University Press, 1991); Robert Goodin (ed.), *The Theory of Institutional Design* (New York, NY: Cambridge University Press, 1999).