Technology in the Culture of Progress

Daniel Sarewitz
Workshop on Transhumanism and the Meanings of Progress
April 24-25, 2008
Arizona State University

Could there be any better mirror into a person's soul than their views on progress? Which litany best encompasses your world view?: the eradication of smallpox, the lifting out of poverty of hundreds of millions of people in South and East Asia, the economic and political integration of dozens of European nations that for centuries were at each other's throats, the defeat of Nazism, Stalinism, and Maoism, the creation of an incredibly egalitarian global information network via the internet? Or are you more comfortable with The Bomb, AIDS, climate change, continuous concentration of global wealth, a billion malnourished people with no access to clean water, and information overload?

What a great subject for academics to argue about from now until eternity. From the inevitability of progress to its impossibility; from its invention as a modern ideal to its persistence throughout history; from its embodiment in scientific truth-seeking and technological advance to its social construction as nothing more than a contextual illusion that justifies particular ways of being and acting, progress can pretty much shoulder any philosophical, cultural, ideological, or statistical burden that we want to place upon it.

At the core of discussions on progress we often find technology, and for good reasons. There is something like an irreversibility to technological change that makes it a particularly inviting frame of reference for thinking about what progress might actually mean. Technology offers a continually, if unevenly, expanding domain of increasing human control and power in the world, and in the process continually transforms the natural and social worlds. Technology instantiates the natural laws that science discovers, and thus represents the worldly application of the scientific search for truth (even if, in many cases, technological application actually precedes and enables the scientific discovery of the laws). Technology also substantiates the imagination of humanity in our inventions, and introduces novelty into the world, capabilities and artifacts that never could have existed before. Technology embodies the Enlightenment ideal of applying rationality to the betterment of humankind.

We have long ago lost our innocence about technology and progress, of course. I don't mean just The Bomb and Agent Orange and big dams, but the idea of technology itself—the idea of a particular artifact or machine that just does its job. We have learned that technologies emerge from social systems, that they reflect and internalize power relations and cultural assumptions. We recognize that social systems are actually techno-social systems; that these systems impose certain orders of behavior on our lives about which we have little choice—who among us is not enslaved by e-mail?; that these systems lock in paths of dependency that make a mockery of human agency—just try decarbonizing the global energy system!; that they enable hierarchies of expertise, influence, and

exploitation—who, anymore, can argue with their auto mechanic? We know that technological systems are now as complex and pervasive and incomprehensible as natural systems; in fact we know that the distinction between technological and natural systems is no longer very meaningful. We know that the dependence of modern market economies on continual growth means that we have to keep inventing and consuming new technologies, whether we really need them or not, and that this process of continual innovation and productivity enhancement leads to apparently unavoidable spasms of severe unemployment and socioeconomic disruption and transformation, along with the wealth creation that seems to have become an underpinning for civil stability.

Technology seems to render more concrete and perplexing—and suitable for academic research and debate—the conceptual ambiguity of the idea of progress. Writing in the early 1930's, Lewis Mumford called this the "ambivalence" of the machine, which, he observed, "is both an instrument of liberation and one of repression. It has economized human energy and it has misdirected it. It has created a wide framework of order and it has produced muddle and chaos. It has nobly served human purposes and it has distorted and denied them."[1]

The latest version of the technology-and-progress Rorschach test is brought to us by an area of innovation lumped under the term "human performance enhancement." Here the idea is that rapidly advancing research in areas such as nanotechnology, information technology, genomics, robotics, and cognitive science are converging to create the potential, in coming decades, for augmentation of human physical and cognitive abilities to levels and at speeds far beyond anything that has come before. The most avid promoters of this area of technology call themselves Transhumanists, to signal the idea that humans are now in the process of transitioning to a new state of advanced species capability. As explained by the Oxford philosopher Nick Bostrom, the Transhumanist "vision, in broad strokes, is to create the opportunity to live much longer and healthier lives, to enhance our memory and other intellectual faculties, to refine our emotional experiences and increase our subjective sense of well-being, and generally to achieve a greater degree of control over our own lives."[2]

So Transhumanism is a bald claim for old-fashioned, technology-induced progress, for things, generally, getting better because of the development and use of technologies, in this case applied directly to making human bodies, genomes, and brains better than they have even been before. And perhaps because we, ourselves, are the central objects of innovation, transhumanism provokes engagement with all the ambiguity and ambivalence of technological progress in spades. This is our actual selves we are now redesigning, or so it is claimed, with ardor on one side, and grave reservations on the other. What's going on is nothing novel at all; humans have always been in the game of transforming themselves with technology. What's going on is radically new, we are now on the threshold of intervening in our own evolution, in the exponential expansion of our own cognitive capacities, in an irreversible blurring of the human-machine boundary. The

---

[1] Mumford, L, *Technics and Civilization*, p. 283, 1928.

[2] http://www.transhumanism.org/index.php/WTA/more/transhumanist-values/

benefits to humans will be wonderful as we achieve new levels of intelligence, creativity, agility, even perhaps wisdom. The threats to humanity are profound, as we alter our natures in ways that may erode the foundations of society, challenging our commitments to justice, democracy, and the very notion of human dignity. And so on.

Now amidst all this ambiguity and ambivalence and hope and fear and philosophical mud wrestling, I am wondering, well, is there anything at all we can say about what is actually likely to happen? I don't mean specific predictions about the technologies themselves; it's fine with me if Ray Kurzweil manages to download his brain into a computer, or if William Bainbridge manages to get his archived personality launched into space, but I sort of doubt that most people harbor such aspirations. In any case, if this human enhancement thing goes anywhere, it will be because people use the technologies, and so that's the way I want to enter into this discussion. And I want to do this by positing the existence of something that no one of any intellectual sophistication could possibly believe in anymore: technology itself. For example, I want to claim that this watch I'm wearing embodies a technological function that I find useful—letting me know what time it is. And in acknowledging this function I'm not at all suggesting that, say, the development of the chronometer in the 18th century did not advance the pace of colonialism by allowing more precise marine navigation, or that my father was not manipulated into buying my particular watch by clever advertising, or even that my need to know the time is not a cultural expectation that assaults my individual autonomy. It's all probably true, and even so my watch helps me navigate my days because it reliably does what I want it to do.

And a related thing I want to do that differs from many discussions about enhancement and Transhumanism is focus on things that are actually happening now, so that the opportunities and dilemmas are not distanced from our sense of the difficulties of the world we live in, but instead are conditioned by a recognizable context, and so they are not subject to distracting debates over technological possibility. In a world already suffused with technology, people will be making choices about *new* technologies based in part on their experience of what actually works and what does not, more than on the unbounded promises made about technologies that don't yet exist. These choices then will help determine how the technologies actually evolve, in ways we cannot possibly know far in advance.

As well, it seems reasonable to expect that, as future dilemmas about human enhancement unfold, they are going to feel pretty much the way that the current ones do, you're not going to wake up one morning and suddenly find yourself in a world where you can buy brain-machine interfaces that will boost your IQ 100 points, or genetic modifications that will render you impervious to aging. To the extent that such ideas—staples of the Transhumanist debate—are even remotely on the horizon, they will be approached slowly, unevenly, with front page claims of amazing advances one day, and page seven revelations of disappointed expectations a year later.

With this in mind, let me start on what I hope is firm ground: Making and using and improving on technologies is, like eating, sleeping, lying, procreating, imagining,

something that people do, it is part of being a person. Mumford wrote: "The age of invention is only another name for the age of man."[3] And part of this innateness is that, at any given moment, humans are utterly, irrevocably, dependent on some set of technologies and technological functions for their survival. Of course the nature of that dependency changes over time, sometimes in radical fits and starts, but it remains, nonetheless, a condition of our condition. And what this means is that assessments of future technologies are relative to the current technological state, not relative to some pre- or non-technological state. Technology is always part of the context.

Second, I want to reinforce this point about functionality. Technologies are meant to do things, to accomplish particular goals or tasks. Technologies are, in effect, cause-and-effect machines, linking a human intent to a particular consequence via the embedded function of the technology, and often doing so with very high reliability, often—but not always—much more than could be achieved using an older technology aimed at the same function, or without the technology at all. And the fact that a technology emerges from a cultural and political choice process, or that a lot of technologies seem to serve no useful purpose other than to advance the profits of the producer, or are designed to do things like kill people, or end up doing things that they were not designed to do, is not problematic for my assertion of functionality.

Now I want to do something academically incorrect and connect technology to a very modest notion of rationality. This modest notion is that when people make a decision, usually their intent is to achieve the objective at which their decision is aimed. I want to suggest that, at the level of individual use, technologies are volition enhancers: they often give us a better chance of accomplishing what it is we want to do than would otherwise be the case. I could have decided to walk or bike from Boston to Tempe to present this talk; I could have planned out my cross-country sojourn with appropriate care; but the likelihood of my arriving in time would have been much less than by doing what I did—flying on one of our infuriatingly unreliable airlines. By condensing as much of the cause-effect involved in my cross-country trip into a single, contained technology that is, in fact, extraordinarily reliable—the jet aircraft itself—I maximized my potential for achieving the intent of my decision: to be here at this moment.

But this raises a third point, that technologies are on the one hand reliable cause-effect machines for our use, but they are also components of complex systems whose connections to the functionality of the individual technology is not necessarily very obvious. Technologies, that is, inhabit two rather independent realities One reality is of the immediate effectiveness of the technology itself as it is used by those trying to accomplish something—the jet airplane that carries me with incredible reliability from A to B. The other reality is of systemic complexity—the air transportation system that seems in many ways to be the absolute embodiment of irrationality and dysfunctionality, whether it's the insane pricing system, the absurd inefficiency of boarding and security processes, the significant inequities in access, the constant delays, the inability to spend

[3] Mumford, p. 60.

one's frequent flier miles, not to mention the almost continual financial insolvency of most airlines.

Now we are back in the comfort zone of ambivalence and ambiguity and mud wrestling. This system complexity that accompanies technologies raises another facet of the question that we all appreciate: the likelihood of unintended consequences. Technologies do not act in isolation, they are connected to other technologies, activities and phenomena that may interact in ways that no one is able to predict or control. So the functionality embodied in the airplane that reliably carried me from Boston to Phoenix is also implicated, for example, in the rapid spread of exotic infectious diseases such as AIDS and SARS, and of course in terrorist attacks that marked a major inflection point in recent U.S. history. Technologies often surprise us because they introduce into society novel capabilities and functionalities whose diverse uses are constantly being expanded and discovered—that's just what we do, remember—and which interact with other technologies, and with natural and social phenomena, in ways that cannot be known in advance.

So here's the sharp contrast I want to bring out, where we as humans by definition live in a world of technology, in cultures of technological innovation, upon which we depend for our survival, and which condition our behavior, and this dependency on one level reflects our efforts to exercise our intent across an increasing breadth of human domains to accomplish certain tasks with incredible reliability, even as it enmeshes us in complex socio-technical systems of Kafka-esque incomprehensibility and capriciousness, systems that themselves demand continual technological updating as part of the process of adjusting to unfolding and unpredictable complexities.

But I want to get back to my airplane. I experience it in two ways: as an extraordinarily reliable piece of technology, and as a component in an extraordinarily irritating transportation system. I remember when I first grew comfortable with the fact of airplane reliability—when I stopped gripping my arm rests during take-offs, landings, and turbulence. It was 1975; I had dropped out of college and was working on a cod-fishing boat based in Boston Harbor. The boat was small and we fished near shore. The cod fishery was already in decline, and we had to switch from using gill nets to tub-trawls—long lines with hundreds of baited hooks on them—in order to get a decent catch. The costs of the bait and fuel were barely covered by the catch, and I was paid in sandwiches and all the cod I could eat. So we were using these pretty elemental technologies, and meanwhile, we were usually within sight of Logan Airport, and planes were continually landing and taking off, every minute or so, with a reliability and precision that made our own elemental struggles to bring in a few fish seem out of a different world.

Our ability to succeed as fishermen depended, of course, on some technologies—the boat, the nets, and when they failed, the tub-trawls, but the most important parts of the system—the fish, and the ecosystem that sustained them—were outside of our control. The jets were part of a technological system that internalized almost everything that was necessary to do for their functioning. The design of the wings was such that the planes had no choice but to take off; the design of the jet engines was of an elegance and

simplicity to provide thousands of hours of reliable operation. Not only that, but the design of the planes was subject to continual, incremental improvement through, for example, the development of lighter, stronger alloys, more precisely machined parts, and more aerodynamic designs. But there was, in contrast, no engineering ourselves out of the absence of fish, in fact better engineering made the problem worse.

Of course this isn't the whole story. It's not just that the planes are reliable—it's that the reliability is completely obvious—and that unreliability is completely unacceptable. People agree on what constitutes the essence of a reliable airplane—one that gets you to your destination without crashing. This is a value that is widely shared. And pretty much everyone agrees that, if you want to travel many miles in the shortest possible time, airplanes are the best option. So there's a sort of transparency to the effectiveness of airplanes that transcends ideological or religious or philosophical differences. Not only that, because the viability of the air transport system—for all its dysfunction and irrationality—depends utterly on the reliability of the airplane itself, the system has to figure out how to maintain that reliability at all levels of operation, from designing, building, and maintaining aircraft, to managing air traffic, to detecting wind shear. The feedbacks from failure, and the incentives to learn and behave appropriately, are very clear and strong. Last year there were more than 10 million departures by U.S. airlines, with no fatal crashes,[4] an incredible record that reflects continual learning and improvement as the size and complexity of the system grows. And yet I recently arrived at an airport only to find that my e-ticket seemed not to exist. Obnoxious persistence on my part led to the discovery that the ticket, which I thought I had purchased while on the phone with an airline employee who, it turns out, worked in Bangalore, had somehow not been entered into the system. So the airlines cannot even manage to reliably sell tickets to the customers who keep them in business. The feedback from my rage provided little opportunity for learning. Yet not only do they manage keep their planes in the sky—more and more of them, in fact—they keep developing improved planes—faster, more energy efficient, more reliable.

So our own experience with flight is filled with contradiction: we are powerless consumers, buffeted by the turbulence of organizational incompetence that seems to grow more pathological with the years, even as we can justifiably maintain a high level of confidence in the operation of the core technology, which itself is subject to pretty continual improvement. Very strange.

Now what can this possibly have to do with Transhumanism and human enhancement? Well, one thing to be said is that there was a time in the past when the idea of millions of humans flying around in the air would have been seen as a considerable enhancement over current capabilities, even if now it is rendered not only mundane but largely aggravating. Given that many technologies aim in some way to create a human ability to do something better than could be done without that technology, one might reasonably say that the whole human commitment to technological innovation could just be re-framed as a commitment to the enhancement of human capabilities. There is certainly

[4] US Bureau of Transportation Statistics website

something to this position, but it is also the sort of irksome argument that some bioethicists and Transhumanism advocates have used to make sweeping claims on behalf of technological promiscuity. On the other hand it does seem to me that the capacity of technologies to allow people to more effectively exercise some modest form of rationality in the world actually gets us closer to the heart of the matter, to the issue of what it actually will mean, on the ground, for people to be "enhanced."

What I particularly want to call attention to is this difference between a level of action where people's intent has a good chance of being translated into a desired outcome, and the larger complex, dysfunctional system where the consequences of action become pretty much impossible to map out, so that projecting intent becomes at best a matter of trial-and-error, and at worse close to futile.

It would perhaps be an understatement to suggest that what I'm saying goes against the grain of much thinking about technology and society for the past 40 years or so. Classic works that helped to define this field, like Jacques Ellul's *The Technological Society*, and Langdon Winner's *Autonomous Technology*, looked at the ways that societies restructure themselves to accommodate evolving technological systems and saw, for the most part, a ceding of human agency and authenticity and democracy to the demands of the system. The functionality of the technologies themselves was a trick, a sleight of hand aimed at distracting us from what was going on at the higher level. And it's true, complex technological systems do make a mockery of the Enlightenment fantasies of rational control over our affairs and dominion over nature. The work of Ellul, Winner, Mumford and others is suffused with an appropriate sense of befuddlement and resignation about what we can actually do to bring the technological systems that we create under our more direct and effective and democratic control.

Yet when, for example, Winner talks about "reverse adaptation—the adjustment of human ends to match the character of [technological] means"[5] to describe our subservience to technological systems, the "compared to what" question kicks in. When was that golden age of harmony and agency and freedom from adjusting our ends to the available means? This is not, I want to emphasize, my apology on behalf of complex technological systems, it is my suggestion that maybe there is someplace else to look at the way that technologies mediate people's lives, not just in the vast, hegemonic, and unthinking system in which technologies are embedded, but at the local, homely, quotidian level of actual use.

I am very taken with vaccines as an illustration of many of the possibilities, conflicts and dilemmas raised by Transhumanism and the technological enhancement of humans, so I want to use them as a way to push a little harder on this contrast between larger system and local use. Vaccines represent precisely the sort of internal, technological enhancement of human biology that seems at the core of the Transhumanist agenda. We are introducing a foreign material into our bodies, stimulating a response by our immune system, and enhancing our resistance to a variety of infectious diseases as a result. In

[5] Winner, L, *Autonomous Technology*, p. 228, 1977.

many ways, in fact, vaccines represent everything one could possibly want in a human enhancement. Extended, and in some cases even life-long, immunity to a variety of diseases can be delivered through a process that takes seconds to administer and which confers its benefits with remarkable consistency and reliability. The vaccine can be administered easily, by people with minimal training, in just about any setting. The effectiveness of most vaccines is quite apparent; most famously, smallpox has been eradicated from the globe, we've come close with polio, and vaccines have played an important role in the reduction of childhood disease and mortality in many parts of the world. While vaccines have often stimulated opposition on both moral grounds and because of concerns over risks, on the whole they are widely accepted and embraced. Perhaps most interestingly, vaccinating against infectious disease has an inherently democratizing aspect to it. The more people who get vaccinated, the better off everyone is. The benefits of herd immunity are a strong stimulus for public policies that encourage widespread and equitable distribution of vaccines. This is a rare case where ndividual enhancement is strongly linked by vaccines to a wider communitarian benefit.

The contrast between the effectiveness of vaccines and the chaos of the system within which they are administered is perhaps even more conspicuous than for air transport. The U.S. health care system, for example, has become an emblem of dysfunction, of inequity, and of immunity to beneficial political reform, and yet most people manage to get necessary vaccinations, and enjoy the protection that they confer. Even in nations with little public health infrastructure, vaccines have often proven to be a powerfully effective intervention for improving health. So vaccines strongly illustrate the dual realities of technological effectiveness and system complexity.

But what makes a vaccine effective relative to other approaches to reduce the toll of infectious diseases? I want to explore this briefly by looking at the case of malaria, a disease for which no vaccine is yet available. One of the more conspicuous failures of modern science and technology policy has been the relative neglect of research on diseases such as malaria that especially afflict poor people living in poor nations. In the past decade or so, private philanthropy, sometimes in partnership with international organizations, national governments and the private sector, has stepped in to try to reduce this imbalance. For example, research on malaria, and malaria vaccines, has increased significantly. Yet the technical challenges to developing a malaria vaccine are daunting, and it is unclear how long it will take for research to yield useful vaccines, or even if they are possible. In the meantime, malaria is killing at least a million people a year, most of them children, most of them in Africa.

In the absence of a vaccine, there are several major cooperative efforts under way to promote prevention strategies centered around the distribution of insecticide impregnated bed-nets, limited indoor spraying of insecticides, and other measures. In many ways impregnated bed-nets are a very appealing technology. Low cost, low tech, simple to use. Small is beautiful. Where impregnated bed-nets are in wide use, the spread of malaria declines rapidly. Yet the first decade of concentrated effort to promote widespread net use in malarial Africa was widely deemed a failure, with incidence of childhood malaria actually increasing. There are many reasons alleged for this failure,

ranging from policy disagreements over appropriate distribution methods to bureaucratic incompetence at organizations like the U.S. Agency for International Development and the cooperative Roll Back Malaria campaign, to simple lack of use by those who received nets.[6] As of 2006, despite significant expenditures and effort, only about three percent of African children in malarial regions were sleeping under impregnated bed-nets.[7]

In response to these disappointments a policy consensus now seems to be emerging around an approach called Integrated Vector Management. IVM combines bed nets with other interventions in a way that, according to the organization Africa Fighting Malaria, "means tailoring a variety of preventative interventions to local context."[8] As described by the World Health Organization, the characteristic features of IVM include "selection of methods based on knowledge of local vector biology . . .rational use of insecticides . . . and good management practices. The specific mix of these interventions depends on local factors, such as the type of mosquito and malaria parasite, climate, cost, and available resources."[9]

Preliminary results of IVM in several countries have been very positive, and have created a sense of optimism about the prospects for making real progress in combating malaria. According to a report by McKinzie and Company commissioned by Roll Back Malaria and released to the public at no less a high-brow setting than the Davos World Economic Forum:

"[A]n investment [in IVM] of approximately $2.2 billion a year for five years . . . can achieve full coverage of prevention and treatment measures in the 30 hardest-hit malaria endemic African countries, which together account for an estimated 90 percent of global malaria deaths and 90 percent of malaria cases in Africa. . . . Over five years, this effort is expected to:

- Save 3.5 million lives
- Prevent 672 million malaria cases
- Free up 427,000 hospital beds in sub-Saharan Africa
- Generate more than $80 billion in increased GDP for Africa."[10]

I have no expertise whatsoever in malaria control, and the alleviation of human suffering entailed in such a radical reduction of malaria would be a magnificent achievement—true progress that we should all hope for. But it is hard to be optimistic about these predictions, for two very different reasons. Let's imagine that, instead of applying Integrated Vector Management to the prevention of malaria, we had a reasonably effective vaccine. What would be different?

---

[6] E.g., Kyama, R. and McNeil, D. Jr., Distribution of Nets Splits Malaria Fighters, *New York Times*, October 9, 2007.

[7] Hill, J., Lines, J., and Rowland, M, Insecticide-Treated Nets, *Advances in Parasitology* 61: 77-126, 2006.

[8] http://www.fightingmalaria.org/issues.aspx?issue=4

[9] http://www.who.int/malaria/integratedvectormanagement.html

[10] http://www.malarianomore.org/businesscase/

Well, certain things might be the same. There would no doubt be controversies over appropriate policies for vaccine delivery, and organizational dysfunction at various levels, and lack of required health care delivery infrastructure and other challenges that have bedeviled efforts to improve public health in parts of Africa. But the essence of IVM, the key to its success, is "tailoring a variety of preventative interventions to local context." Yet surely the less tailoring that is necessary to achieve a desired result, the more likely success will be. To the extent that IVM is dependent on responding to local context, it is also dependent on managing and implementing knowledge and action in ways that are particular to that context. In any particular setting, the appropriate mix of interventions— "a combination of bed nets, medication, spraying, and environmental management"—must be determined, and the institutions and people responsible for delivering the interventions must act appropriately. IVM, that is, is a complex combination of activities, activities that require organizations to behave in particular ways in particular settings, and, crucially, where no single activity embodies the essence of malaria prevention. In these sorts of complex organizational settings, learning is particularly difficult because it is often not obvious how the lessons from one context might apply to another, and disagreement over competing approaches is common because the links from cause to effect are difficult to fully specify—numerous interpretations of success or failure may be plausible.

The key selling-point of IVM—its sensitivity to context—is nevertheless also its weakness. It's not that current approaches to controlling malaria should not be tailored to context—surely they must be—but tailoring to context is hard to do. A really effective intervention is one that renders context irrelevant. If a reasonably reliable malaria vaccine were developed no doubt there would be challenges related to cost, manufacturing, and social acceptance of the vaccine, but much more of the intervention itself—the action that leads to the desired outcome—is embodied in the technology itself. What you have to do to succeed is clear and always the same: vaccinate people. If people get vaccinated, they will with high reliability be immune, regardless of where they live, how they live, what they believe. A person may or may not reliably use the bed-net they are given, but once they have been vaccinated the problem is solved—the locus of reliability shifts from the individual to the technology. The process of delivering the vaccination is what we might call a "shop-floor" activity, where effectiveness depends little if at all on the larger organization setting. As with airplanes in the dysfunctional transportation system, most of the cause-and-effect directly relevant to solving the problem is localized in one particular technology whose performance can be easily measured, where performance is largely insensitive to the surrounding institutional context, and where pretty much everyone agrees on what counts as success.

Solving a problem is hard when you don't have a way to condense the key cause-effect elements of the solution into a particular technology or routine that can be administered at the shop floor. Perhaps there will never be effective malaria vaccines, in which case IVM may well be the best path to reducing malaria in poor countries, but it is likely to be a very indirect path, one that does not always reach its goal. An effective vaccine would do the job better.

Of course many important problems can't be technologically embodied in this way. The report I mentioned earlier on IVM predicts not only that a $10 billion investment will lead to reduced malaria cases and mortality, but also to enhanced wealth creation--$80 billion over five years, to be exact. One should expect that healthier people are better able to contribute productively to an economy, but there are so many other confounding factors, from education levels to environmental conditions to the quality of governance to the global trade situation that any prediction based on links between changes in malaria incidence and changes in wealth creation is at best a vague and hopeful guess about complex system behavior. This problem is nicely illustrated by the work of anthropologist Peter Brown,[11] who tested the hypothesis that malaria was blocking economic development on the island of Sardinia in the period after World War II. He concluded that the "macroparasitism" of landowners drained thirty percent of the production capacity from peasants in the form of rents, while the "microparasitism" of malaria accounted for less than 10 percent reduction in their gross production. And here we shouldn't expect a vaccine to do any better than bed nets; now the goal—creating wealth—cannot be captured and internalized by a particular technology. In fact, if creating wealth is your goal, there could be much better routes to progress, for example changing patterns of land tenure, or improving levels of education. But of course these goals are themselves very hard to make progress on.

Let me push a little farther into enhancement land. Cochlear implants are electronic devices that provide deaf people with a sense of hearing by direct stimulation of auditory nerves. Unlike hearing aids, which just amplify sound, implants can give profoundly deaf and severely hearing-impaired people the ability to sense and interpret sound, including speech. First tested in the early 1960s, cochlear implants were approved for use in the U.S. by the Food and Drug Administration in 1985. Today, on the order of 100,000 devices have been implanted worldwide, most of these in the past decade.

Opposition to cochlear implants has been centered around the idea that deafness is a reflection of human diversity, not a deficiency to be corrected. From this perspective, deaf culture is as rich and worthy of protection as any other distinctive culture, and deaf people don't need to be "enhanced." Using sign language, deaf people communicate richly with each other and with hearing people who also sign. Obstacles to complete and equal participation of deaf people in a society dominated by hearing people is a reflection of the structure of society, of institutionalized injustice, not of the attributes of deaf people. The appropriate ethical goals, therefore, are to eliminate obstacles to full participation in society by deaf people in society, and to adopt a stance of complete acceptance of deaf culture. Cochlear implants represent a threat to these goals, and to the sustainability of deaf culture. They should therefore to be resisted. In support of this position, at least one deaf couple has sought out a sperm donor with a long history of family deafness to ensure that their child was also deaf.

---

[11] Brown, P.. Microparasites and Macroparasites, *Cultural anthropology*, 2(1), 155-171, 1987.

The reasons why people choose cochlear implants are apparent. The main ethical subtlety lies in the fact that many devices are implanted in young children who cannot consent to the operation or to the enhancement—but this is hardly a unique situation—and that in such cases it is presumably hearing people for the most part, rather than deaf people, who are making the decision. Now my point here is not to disparage opposition to cochlear implants; in fact, the position is based on a generous vision of social justice and equity that I support. But justice and equity can be served in different ways. If the goal is to create a society where deaf people have all the opportunities and benefits of hearing people, and two paths are open—the struggle for the complete and equal rights and access of those without hearing, and the widespread adoption of cochlear implants—it is likely that one path will be much more difficult, uncertain, frustrating, and protracted than the other. As in the case of malaria, one option pushes the complexity of the larger system to the background by embodying most of the relevant cause and effect of the immediate problem, and in so doing it radically reduces the political and organizational challenges involved in making progress to achieving the goal.

My alert listener has perhaps sensed that, at this point, we begin to get into a little trouble on the ends-means front. Is the goal to ensure that deaf people can participate as fully and equally as possible in society? Or is it to ensure that deaf culture is sustained and fully accepted and included in a vibrantly diverse society? It could be true that a society that fully embraced deaf culture as an expression of human diversity would be a more just and equitable society than one which radically reduced deafness through the widespread use of cochlear implants. But because cochlear implants exist, modifying deaf individuals technologically so that they can participate in hearing culture is simply an easier task—a much easier task—than stimulating the political and behavioral change necessary to fully include, nurture, and sustain deaf culture and in the process make implants irrelevant. The dilemma exemplifies Langdon Winner's complaint about adjusting our ends to match the character of our technological means. But this adjustment also allows us to act more effectively. I am not making an ethical argument for or against cochlear implants, I am making an observation about the likely success of different paths of social change, and a weak prediction about the choices that people on the whole—exercising my version of modest rationality—are therefore likely to make.

So this is a tough dilemma that pits the ease of a reliable technological intervention against the hard and unpredictable slog of political struggle. What gives the technological option the key political advantage is its functionality, its effectiveness—if it didn't actually do what it claimed to do, its competitive advantage against the non-technological option would be harder to sustain. And so the effectiveness of the implant itself has a political meaning, a built in political impetus, and the implant becomes an attractor of various constituencies that want to get a certain thing done, because it does so more reliably than other avenues to a related end. The dilemma is particularly disturbing because it rubs our noses in the possibility that what seems like right action—the quest for a more tolerant and inclusive society—may in the end be a less beneficial path to follow than the use of the technological intervention—not because it wouldn't be better

to have a society that didn't need the technological fix, but because the effectiveness of the fix is so much more reliable and predictable than the political path toward progress.

I want to pursue this tension one uncomfortable step farther. One of the most conspicuous sites for technological enhancement of humans is the process of childbirth. The influence of technology on the birthing process has been pervasive and profound and will likely become more so. One may also feel it has been alienating and dehumanizing. At the same time, the industrialization of childbirth, through the application of technologies like labor-inducing drugs and heart monitors, through standardized procedures like the C-section and the Apgar score, has also transformed childbirth into an extraordinarily reliable and predictable process relative to what it once was, where, in rich countries, infant mortality at childbirth has declined from several hundred per thousand in the 18th century to ten or less per thousand today, and maternal mortality has declined from as high as ten percent or so as recently as the late nineteenth century to less than one in 10,000 today.[12] When you combine these trends with the rise of assisted reproductive technologies like in vitro fertilization and the increasing ability to nurture radically premature babies to term, one can imagine that we may be on a course toward pregnancies that are completely technologically mediated, perhaps even outside the womb, to deliver, as it were, utter reliability in the outcomes of childbirth.

This is a trajectory whose wonderful benefits are undeniable even as the continual intrusion of technology on pregnancy and childbirth may reasonably offend our sense of what is appropriately human. This offense may seem to be magnified when we think about a related problem: the entrenched inequities in birth outcomes in the United States. Infant mortality among African Americans, for example, is on the order of twice what it is among whites; and the overall rates of U.S. infant mortality have long been unconscionably high relative to other rich countries, mirroring America's greater levels of socioeconomic disparity. So we pursue all this technological change in our affluent society but meanwhile we can't even do what is necessary to ensure that poor and minority babies have the same chance of surviving as white and well-to-do babies.

But there are, it turns out, two twists to this tale. First, over the last few decades, infant mortality rates among poor and minority babies in the U.S. have actually declined pretty much at the same rate as declines among the babies of more well-to-do parents. So, while the disparities remain distressingly resistant to change, the absolute outcomes have improved for everyone. These declines are apparently explained almost entirely by technologies; shop-floor interventions, in the delivery room, that have provided tremendous benefits to poor and well-off alike.

The second twist is that there have been substantial efforts to address the inequity problem, and they have largely failed. More than forty years of policies aimed at increasing the quality of prenatal and maternal health care and nutrition among poor women in the U.S. through Medicaid and other programs have had little or no positive effect on birth outcomes nationwide. These worthy efforts turn out not to have narrowed

[12] e.g., Porter, R., *The Greatest Benefit to Mankind*, 1997.

the mortality disparities.[13] The specific reasons for lack of progress are of course debated among the experts. The causes of high infant mortality rates among poor people are complex, and deeply embedded in broader problems of socioeconomic inequity that continue to resist political solution and effective policy intervention.

Obviously I am not arguing against engaging in the hard political battle for greater socioeconomic equity in our society; the growing concentration of wealth in our already wealthy nation and world is a continuing moral affront. Rather, my point again is that when the essence of a problem is amenable to capture by a technological intervention, real progress can sometimes be made very rapidly, whereas political paths to solving a bigger, underlying problem are likely to be much slower and less satisfactory. This is what we are seeing in the infant mortality case.

Yet somehow the technological path may seem less ethically satisfactory than the political path, because it leaves intact the underlying social failures that contribute to the inequity. Again, we adjust our ends to suit the available means, and this may create some reasonable sense that the technological path provides us with an excuse for not taking the political path—that the available means are distracting us from the more important end, from doing what is right, which is to solve the problem by making society better, by reducing inequality, rather than by isolating the problem from its social context through a technological fix.

In a talk I gave several nights ago I tried to make the point that technological enhancement of humans is not going to help us confront the most fundamental political challenges faced by our societies, challenges underlain by a combination of value conflict and uncertainty about the future consequences of our actions. The stubbornly persistent socioeconomic inequities that continue to fester in the United States, and that underlie our poor performance on infant mortality, are one example of this sort of problem that is not amenable to technological solution. Here I am trying to make a complementary point: that in reality there is no easy path to addressing such fundamentally political challenges, but that technologies can sometimes help find a shortcut to dealing with some of the particular consequences of these challenges.

So I don't buy into the framing of Transhumanism as offering a choice between two alternative technological futures, one essentially utopian, the other essentially dystopian. I want to suggest instead that there is a scale of experience where one does not have to give up one's sophistication about the complexity of the world to accept the possibility of modest yet encouraging technological progress. This progress derives from our innate capacity and apparent compulsion as a species to technologically innovate—to take

[13] Lantz, P., Shultz, C., Sieffert, K., Lori, J., and Ransom, S., The Impact of Expanded Models of Prenatal Care on Birth Outcomes: A Critical Review of the Literature (ms. in prep.); Gortmaker, S., and Wise, P., The First Injustice: Socioeconomic Disparities, Health Services Technology, and Infant Mortality, *Annual Review of Sociology* 23: 147-170, 1997.

certain types of problems and capture much of what is difficult about them in physical artifacts that allow us to get around those difficulties.

In making this kind of progress we are creating a domain of increasing control related to solving a particular problem, even as we are also feeding into the complexity of socio-technical systems whose comprehensibility continually slips from our grasp, and often presents us with new types of problems. This seems to me like a fundamental dilemma of the human condition, one that requires continual, balanced attentiveness. Technology is no more the cure for politics than politics is for technology; they need each other, can benefit from each other, evolve with each other, and we are forever stuck with both. Yet I do think we need to recognize and appreciate that there is something different, and special, about technology that under certain circumstances allows us to act in the world with much greater effectiveness than we could otherwise do. In fact I'd go so far as to say that the greatest source of reliable action in human affairs is not our institutions, cultures, or norms, but our inventions. Any approach to solving the many vexing challenges that face the world today needs to accommodate this fundamental, if uncomfortable, reality.