

18

Daniel Sarewitz

Director Consortium for Science, Policy & Outcomes

Arizona State University
USA

1. Why were you initially drawn to philosophical issues concerning technology?

Actually I hadn't thought about this before, but I suppose I am, in fact, "drawn to philosophical issues concerning technology." "When" is easier than "why." In sixth grade—this would have been 1966 or '67—I was the only kid in my class willing to take the "against" position in the class debate about the Apollo mission. I made, as I recall, the opportunity cost argument—why spend all those billions sending men to the moon when there are people without enough to eat in America's inner cities and rural poverty pockets? So here, perhaps, is also a clue about "why." Injustice on all scales has always given me a pit in my stomach, and apparently at the age of eleven I located in technology a key site for the production of injustice.

Several years later, on the other hand, Kubrick's visionary 2001: *A Space Odyssey* opened up for me the grandeur, magnetism, and mystery of space exploration. Perhaps going to the moon and curing poverty were both worth doing. It wasn't until embarrassingly recently, however, that I understood the most important reason why we managed to go to the moon in the 1960s while simultaneously failing to cure poverty—despite the resources devoted to the creation of "The Great Society" in America during that time: The former was easy, and the latter was very hard.

2. What does your work reveal about technology that other academics, citizens, or engineers typically fail to appreciate?

First: The dominant role of technology in the lives of people living in affluent societies is to prevent the collapse of the hedonic treadmill that powers market economies.

It seems to be true that continual economic growth is necessary (yes, yes, but obviously insufficient) for continued civil stability in modern market democracies. It is equally the case that continual technological advance is necessary for economic growth. Furthermore, in my view the most profound empirical finding (and one of the most robust, as well) to emerge from the social sciences in recent decades is that subjective quality of life—people’s assessments of their own level of well-being, of their happiness—is decoupled from economic growth. Almost all nations for which there are data show incredibly stable levels of subjective well being over the past forty years or so, a period of huge wealth creation: increased societal affluence does not translate into increased feelings of well-being.¹ From one perspective, this is a consoling result: more money really can’t buy more happiness! (did we need decades of survey work to prove this?). But from another perspective the detachment of happiness from wealth creation illuminates a troubling incoherence in the relationship between humans and their technology.

Economists of technology like to observe that technological innovation is the dynamo of economic growth, but what they really mean is that the *consumption* of the products of innovation accounts for this growth. Yet when people consume technologies, they are not doing so to make their little contribution to wealth creation, they are doing so to satisfy their own wants (whatever the source of those wants may be).

Indeed, we talk of technology, at its best, as the source of solutions to problems, and we talk of the solution to problems as a key to improving our lives. The converse is also true: because our consumption of technologies leads to the continual remaking of

¹The reasons are dauntingly obvious: happiness flows from the quality of one’s social relations, from deep satisfactions of meaningful work and engaging play. There is no obvious link between more affluence and more of any of these sources of well being. Subjective well-being also flows from a sense of doing well relative to others; more affluence may satisfy this sense for individuals, but not across society as a whole.

the appearance, mechanics, and dynamics of our daily existence, from the way we enjoy music to the way we eat and work and have sex, to imagine ourselves in a previous world is to imagine ourselves in an uncomfortably different world, a world in which most people, presumably, would not wish to live. Give up my i-pod or my portable defibrillator? Forget about it. Yet the data tell us that we don't actually like living in this current version of the world any more than we liked living in previous ones. So our behavior as avid and habitual consumers of wealth-creating and world-transforming technology only makes sense if viewed as society-wide neurosis: apparently we continue to believe that future consumption actually will lead to a better quality of life (or stave off a decline), even though we have never experienced it that way before. Our happiness, satisfaction, and wellness, that is, are coupled to our technological condition only via our vain ambitions for a more satisfying future that never arrives. But if we gave up on technological consumption as the road to a better life then the centrifugal force that holds modern market economies together would dissipate. We are gerbils on an exercise wheel, and our continual running in place is precisely what sustains modern gerbeldom.

Second: Technology is a powerful political organizer.

I take it that the essence of a technology is its internalization of one or more cause-and-effect relations and its ability to reproduce those relations with high reliability, often in a wide variety of contexts. A light switch, a thermonuclear device, scuba gear, a vaccine, or the personal computer and software with which I now write: they are remarkably dependable, the embodiments of reliable action. They are also, apparently, considerable improvements over prior technique, e.g., respectively: candles, conventional bombs, holding your breath, quarantine, typewriters.

A lot of writing and thinking about technology in the post-industrial world has focused on its contested nature, canonically displayed in the debates over genetically engineered foods (GMOs) and nuclear power, and, today, human cognitive enhancement. We haven't thought enough about another aspect of technology: its ability to bring diverse political and institutional players into a well-coordinated solar system of collaborative action. If you think about a relatively uncontroversial technology, childhood vaccines for example, the coordination of social activity that has taken place to ensure their effectiveness is quite extraordinary. The vaccine industry, medical practitioners, health insurers, government

regulators, school systems, local governments, and individual families have all recognized a common interest that allows them to work together in reasonable harmony to ensure that most children in the U.S. get vaccinated. The gravitational center of this near-miraculous degree of cooperation is the core technology itself—vaccines that yield reliable, and desirable, outcomes, and thus motivate the accretion of the necessary network of actors and institutions, which themselves represent an enormous diversity of interests, values, and ways of knowing.

I'm increasingly drawn to the idea that successful technologies can be understood as highly effective agents of political organization. Because technologies are, in effect, reliable cause-and-effect machines, those whose interests and goals are advanced by what a technology does would be foolish not to organize around the use of that technology. Of course those whose interests and goals might be threatened by a given technology may organize against it, and ultimately be disenfranchised by it, but such opposition labors at a huge and inherent disadvantage because the interests of their side are not bolstered by the effectiveness and high reliability of the technology itself. The common image of Luddites is not of fighters for social justice, which they were, but of quixotic, at best, resisters of the inexorable tide of innovation. Today, opposition to the planting and consumption of GMOs, while currently holding a beachhead in western Europe and Africa, is ultimately doomed because of the highly reliable behavior and delivery of benefits that GMOs offer to their supporters. This high reliability in both performance and delivery of benefits is in strong contrast to the variety of interests that need to be held together by a more abstract commitment to opposing GMOs. Technology stacks the political deck in favor of its own side.

I need to be clear about two aspects of this argument that are likely to be misinterpreted. First, I am not saying that there may not be very good reasons for opposing various technologies because of the potential for social or environmental impacts that adversely affect legitimate interests and important values. I'm only saying that if such opposition is framed in terms of using or not using a technology, those who believe that the technology should not be used are shouldering a grave political and practical disadvantage. Political because, as I am arguing, effective technologies are a much more coherent nucleus around which to organize interests than are abstractions. Practical because effective technologies are, in general, much better at achieving the outcomes for which they

are designed than are approaches to problem-solving that do not have, at their core, an effective technology but rather depend, say, on humans behaving or organizing in certain ways. Thus, while I think there is a reasonable basis for concern about GMOs from cultural and ecological perspectives, if the goal is to maximize food production and minimize environmental damage from chemical inputs and land use, then GMOs are likely, on the whole, to define a more successful path than approaches that depend for success on complex, integrated changes in the behavior of diverse actors in the global agricultural system.

Nor am I making an argument on behalf of technological determinism. We all know that technologies and the complex social and institutional networks that grow up around them are products of social systems and human choices. Nevertheless, to fail to recognize the action essence of technology—it's reliable embodiment of cause-and-effect relations, of effectiveness in the world—is to fail to understand the power of technology, in turn, to influence those systems and choices. One might object that the notion of "technological effectiveness" is incoherent because such effectiveness does not exist in a vacuum but demands complicit organization of social systems. But powerful positive feedbacks emerge between such systems and the technologies at their core. If one chooses to alienate one's self from the technologies, then one also isolates one's self from the feedbacks and the powerfully effective actions they enable.

So I've made two points here about technology that I think are insufficiently incorporated into thinking about how society works, and in particular into discussions about technology and positive social change. These points may seem to be contradictory, or at least in tension: first, that technological change in market democracies is a Sisyphean project when it comes to improving people's own sense of well-being, and second, that technologies themselves are powerful engines of social action and, therefore, political organization. These points may appear to be contradictory because it seems sensible that improved life satisfaction would naturally be one measure of the worthwhileness of social action. That is, if technologies embody effectiveness at achieving the ends desired by those who deploy them, wouldn't one expect that they would also add to the satisfaction of those people as well, through the process of fulfilling desires? So one way of interpreting things is to say that the contradiction is real, that technology obviously is at best only weakly coupled to the achievement of life satisfaction,

and therefore that we need to look to other, non-technological sources of public and private value in the search for a path to greater and more meaningful well being. I take this as a pretty widely accepted perspective among political progressives in an era where the threats, failures, and uncertainties of technology are manifest. I am hugely sympathetic with it. But it raises a very troubling problem: if progressives don't take technology seriously, then those that do will have an immense operational and political advantage in the world.

3. What, if any, practical and/or social-political obligations follow from studying technology from a philosophical perspective?

So here's an example from the real world: an academic biomedical researcher with an economic stake in a particular genetic therapy wants to test it on a patient. A bioethicist at the same university is consulted and approves of the experiment (I presume, perhaps wrongly, that bioethicists are supposed to engage in philosophical reflection?). Patient dies, but the scientist has fulfilled the obligation that follows from studying technology from a philosophical perspective, right?

At the opposite end of the spectrum we have the precautionary principle, two words which I have a really hard time typing out or saying without first appending the disclaimer "so-called," and which is probably the crowning jewel (or maybe the only jewel) of philosophical "obligations" relating to the governance of technology. Of course I understand that in its generic version the PP just means that given uncertainties about the health and environmental risks of new technologies, erring on the side of prudence is a good thing, as is providing a regulatory and psychological counterbalance to the exuberance of profit-motivated technology producers. Who could argue with this, until you try to really figure out what "erring" and "prudence" mean when you're talking about complex, networked, technological systems in which cause-effect relations are pretty much impossible to parse. For example, how easy is it to know what the truly precautionary position is for using DDT against malaria in Africa?

I've argued, in essence, that rejecting technology is rejecting power. If you're serious about your politics, then this is a bad place to begin. It seems odd to me that progressive politics has been, on the whole, willing to adopt science as a legitimating source of cultural power while maintaining a high degree of skepticism about

technology. Science is inherently contestable and open-ended, and thus equally available to all parties. It is an ineffectual tool for political persuasion. It also does not translate easily into action. Technology, on the other hand, can short-circuit political conflict not by persuading diverse interests and values to converge, but by allowing them to co-exist in a shared sense of practical benefits derived from the action embodied in technology. A great example is the story of how the nations of the world came to an agreement in the late 1980s to phase out the production of chlorofluorocarbons—a technologically and economically important refrigerant and solvent that also happened to destroy the stratospheric ozone layer that protects Earth from harmful ultraviolet radiation.

Popular narrative has it that the CFC-ozone problem was solved when the science proved that CFCs had caused the Antarctic ozone hole, and the world, faced with definitive knowledge, took effective action. This is a story of science forcing right behavior, but it is bizarrely incomplete, because it ignores the question of what “effective action” actually means, and how it is to be taken. The many nations that agreed to phase out CFC production did not, in so doing, decide to live without the benefits of keeping refrigerators cold. The missing element here is the technological. Neither the politics nor the science surrounding the ozone debate was settled until Dupont invented a CFC substitute, hydrochlorofluorocarbons, that was less destructive of ozone. The availability of an alternative to CFCs made it possible at once to meet the goals of three distinct constituencies: those whose primary interest was to protect the ozone layer, those whose primary interest was to make money producing chemicals, and those, especially in the developing world, who were unable to give up on the benefits that CFCs alone could provide in an economically viable way. This story is perhaps less satisfying than the tale of science convincing people to make sacrifices for the good of the planet and humanity, to do the right thing regardless of worldly consequences, but it has the virtue of actually explaining how effective action was able to come about. Sometimes a technological fix is a beautiful thing.

A predisposition toward rejecting technological approaches to social problems—PP is rooted in the predisposition—fosters impotence in two ways, first by rejecting the worldly effectiveness of technology, and, second, by doing without the political organizing value created by such effectiveness. This phenomenon is on depressingly stark display with the problem of climate change,

where the political agenda has focused on using science to compel political action, with the result that we've got lots of science combined with a political agenda that has had no discernible effect on preventing either climate change or the damaging impacts of such change. Attributing this failure to lack of "political will" simply restates the problem, since "lack of political will" is just a synonym for "not wanting to do things my way." Climate change will have many solutions, but the transition to a carbon-free energy economy must largely be a technological one, so if progressives are really serious about climate change, then technological systems should be central to their political agenda. Why hasn't it been? I find it really difficult not to think that the political framing of the climate change problem is largely impelled by concern about the moral meanings of affluence (excessive consumption and the like; return to question 2, part 1, above) rather than by concern about the consequences of climate change itself. Take this test: if in the next decade or two technologies for capturing and storing carbon emissions create the realistic prospect of essentially decarbonizing the global energy system, would you be satisfied, or would you feel somehow cheated of the opportunity to make people work harder for their redemption? Where does obligation most potently lie? On the side of trying to get people to behave differently, or of trying to advance technologies that reduce the need for behavioral change?

4. If the history of ideas were to be narrated in such a way as to emphasize technological issues, how would that narrative differ from traditional accounts?

I know very little about the history of ideas, and even less about what constitutes a "traditional" account. Given received myths about inexorable and inevitable progress in the world (and their antithetic myths of the coming eco-socio-techno-dystopia), it's not clear to me that the problem is lack of emphasis on technology but lack of clarity about how our philosophies and our normative commitments relate to our technological narratives.

5. With respect to present and future inquiry, how can the most important philosophical problems concerning technology be identified and explored?

In the past year I've heard compelling philosophical treatments of the moral obligation to rapidly pursue human enhancement technologies and the obligation to abandon such pursuit. I've heard a neuroscientist, fighting back after a series of philosophical and political assaults on enhancement research, ask why science aimed at helping people overcome their cognitive disabilities and limits should be approached from a precautionary perspective but not energy research aimed at helping the world overcome its dependence on hydrocarbons? It's a perfectly reasonable question. I don't know much about philosophy, as is probably obvious, but my observation is that the ethics of justifying biomedical technology and the ratiocinations underlying the precautionary principle, while rooted in opposite assumptions about technology (one permissive, the other conservative), each address a decontextualized and therefore non-existent world where individual agency and artifacts are the proper unit of analysis for addressing the question of what we ought or ought not to do regarding questions of technology. Such approaches, on both the optimistic and pessimistic sides, seem, to me, to fail to engage the fact that innovation is as defining a human attribute as just about anything else you can name, and that the human prospect just cannot be extricated from the technological systems that humans have created and on which we entirely, now, depend.

If technological innovation is understood as a core activity of the human species, as organic as composing music or falling in love—which it is—rather than an elective hobby that can either be pursued or not, then the core question about technology becomes one of governing, of modulating, the innovation activity itself. We ought, it seems, to be striving for a philosophy of the innovation process, of the role of technology system evolution in human affairs, rather than a philosophy of the creation and meanings of technological artifacts. We know a lot about how innovation occurs, especially its decentralization among institutions and sectors and actors. It's a networked process, it's usually incremental, it's not cleanly separable from other social activities and institutions. Human choice is present at many venues, some of which have greater leverage than others, but none of which are hegemonic, in that there is no site that can visualize a future technological system and make it happen according to plan. People who think

a lot about technology and society understand all this quite well, and the key problem of intervention that it entails: the better we can see and understand a technological system, the more difficult it is to have much impact on its evolution, both because of the inertia, or lock-in, that gets built into an operating system (e.g., the problem with automobiles is not just internal combustion engines, it's also roads and filling stations and suburbs, the organization of manufacturing, etc. etc.), and because of the increasing economic and political power of those whose interests are served by the technology. But at earlier stages of evolution, the technologies and interests surrounding them are insufficiently formed to motivate action or to know with any confidence how the systems will evolve.

From a political and policy perspective, the solution to this dilemma is now more-or-less clear. The innovation process needs continually to be subjected to the rigors of pluralistic reflection and discourse. This process needs to take place throughout the complex, disseminated venues where decisions related to innovation are made: laboratories, government research agencies, legislatures, corporate board rooms, retail stores. The contexts from which innovation arise need to be as self-conscious, argumentative, and openly normative as any other important social activity with world-transforming potential.

For me, the most compelling and worrisome technologies on the horizon now are those having to do with the enhancement of human cognitive function. Advances in fields ranging from neuroscience to nanotechnology to genomics are creating the potential to increase memory, expand sensory and information processing capabilities, and even enable direct, remote communication between brains and machines, and between one person's brain and another's. Technical experts talk confidently about the potential to enhance intelligence, and also raise the possibility for improvement of intangible qualities like judgment and wisdom—and even subjective well-being, or happiness.

There are plenty of good reasons to squirm at these prospects: the potential for using such technologies to diminish the cognition of our enemies; the injustice and inequity that would flow from inevitable problems of unequal access; the fact that some of the attributes said to be enhanceable, like intelligence, are in part socially constructed and thus internalize existing social power structures; the profound challenge that such technologies could create for the functioning of our fundamental democratic institutions,

and even for economic markets, which have arisen, organically as it were, from a species with a particular range and distribution of cognitive assets.

But an approach to cognitive enhancement that takes rejection as a starting point is an untenable foundation for action in the real world, and is instead a recipe for marginalization. Because the boundary between therapy for those who suffer from cognitive impairment, and enhancement for the “normal,” is fuzzy, the moral case for complete rejection is a difficult one to make. The practical case is impossible to make: in an increasingly globalized technological innovation system, the possibility for suppression of effective technologies that can advance powerful interests is nil, as we see in the continued proliferation of nuclear weapons and GMOs. And if some of these enhancement technologies prove effective, they will instantly attract constituencies who, in benefiting directly from the enhancements, will thus be better able to act more effectively in the world. Transhumanism, the small but loud cultural movement rooted in the idea that part of what makes us human is our continual pursuit of “better” human performance through technology—I enhance, therefore I am—is a force to be reckoned with not because of its tortured philosophizing but because it aligns its own interests with technological innovation.

To bring things to an end, I must confess to a smidgeon of optimism here. It actually does seem to me that public discussions surrounding emerging areas of world-changing innovation, such as nanotechnology, genomics, and cognitive enhancement, are more intense, and are penetrating more deeply and more pervasively into the innovation process, at an earlier stage, than has ever been the case. The site of philosophical engagement here needs to move away from the specific technologies (most of which don’t yet exist and at least some of which may never prove feasible) to the dispersed components of the innovation processes from which complex technological systems will emerge. Innovation is a political activity, which means that there is no optimal outcome. Agency is not really located in individuals, so standard moral yardsticks for human behavior are not particularly relevant. If the principles upon which a good society must be built—justice, tolerance, equity, sustainability, etc.—are applied only as the measures of technologies themselves, then they will be applied impotently, and too late. Rather, they must inform the initial conditions of human choice from which innovation, in all its diversity and unpredictability, can proceed.

Additional Readings

Sarewitz, D. 2005. This Won't Hurt a Bit: Assessing and Governing Rapidly Advancing Technologies in a Democracy, in: *The Future of Technology Assessment* (Washington, DC: Woodrow Wilson Center), pp. 14-20.

Sarewitz, D., G. Foladori, N. Invernizzi, and M. Garfinkel 2004. Science Policy in its Social Context, *Philosophy Today*, v. 48(5): 67-83.

Sarewitz, D. 2004. How Science Makes Environmental Controversies Worse, *Environmental Science and Policy* v. 7 pp. 385-403

Sarewitz, D. 2003. Science and Happiness, in: A. Lightman, D. Sarewitz and C. Desser (eds.). *Living with the Genie: Essays on Technology and the Quest for Human Mastery*, Washington, DC: Island Press, pp. 181-200.

Sarewitz, D. and E. Woodhouse 2003. Small is Powerful, in: A. Lightman, D. Sarewitz and C. Desser (eds.). *Living with the Genie: Essays on Technology and the Quest for Human Mastery*, Washington, DC: Island Press, pp. 63-83.

Guston, D.H. and D. Sarewitz 2002. Real-Time Technology Assessment, *Technology in Culture*, v. 24, pp. 93-109.

Sarewitz, D. 1996. *Frontiers of Illusion: Science, Technology and Politics of Progress*. Philadelphia: Temple University Press.