

All We Lack is the Political Will: Technology and Effectiveness in Human Affairs

Daniel Sarewitz
Arizona State University

Science, Technology, and Public Policy Lecture Series
University of Michigan
November 12, 2007

The world, it almost goes without saying, is a mess. Or perhaps better put, it is a bunch of messes. The rise of global terrorism, the proliferation of dangerous weapons, the continued spread of AIDS, the oligarchy in Russia, the endless conflict between Israel and Palestine, the decline of global fisheries and tropical forests, choose your favorite mess. My own interest has been in messes that implicate, in one way or another, science and technology, for example climate change, a depressing mess if there ever was one.

One of the most discouraging aspects of this particular mess, and of other environmental challenges, is that our scientific understanding is fairly strong—we can see what the problem is, but we seem unable to move effectively and decisively toward solutions. This leads to a sense of huge frustration: we *know* what the problem is, right? If only we had the political will to do the right thing! If only we had the right leadership! If only people would behave differently! Then we could move toward a solution. Of course, such sentiments are really just restatements of the problem itself. To say we *know* what the problem is but we can't solve it because of, say, partisan politics, or vested interests, is just to move the location of messiness from, say, the coupled ocean-atmosphere system to the coupled ocean-atmosphere-society system. Not only that, knowing “what the problem is,” even within the natural system, and knowing how to intervene in that system to create a particular consequence are two entirely different things as well.

Now people who try to think seriously about how to intervene effectively in one mess or another typically start with the insight that we are dealing with complex systems, where cause-effect relations are poorly understood, where contingencies and non-linearities abound, where accurate predictions of future states are impossible, and where the role of the individual is practically buffered into non-existence. Such people, and I am one of them, take seriously the H.L. Mencken aphorism that “for every complex problem there is a solution that is simple, neat, and wrong.” The intervention challenge, it appears, is not to look for easy solutions, but to find ways to manage complexity and indeterminacy. The best we can do is to muddle through; indeed, to try to do better is often to do worse. Muddling through is guided by a strong sense of long-term goals, but an equally strong agnosticism about what the best paths to those goals might be. People who take complexity seriously have sought to formalize muddling through via decision-making strategies with names such as incrementalism, adaptive management, intelligent trial-and-error, and reflexive governance. In the sense that we appreciate the daunting challenges that complexity creates for purposive human action—for making the world a better place—an appreciation of complexity and contingency makes us conservatives, in the traditional and philosophical, rather than ideological meaning of the word. We are

cautious about the prospects for stimulating rapid change toward particular desired goals. The astounding mess in Iraq is only the latest object lesson in what happens when you ignore the wisdom of such conservatism. I think we see similar evidence for having ignored the implications of complexity in the current state of the climate change mess.

A central feature of this way of looking at the world, a feature almost too obvious to mention, is that context matters greatly. Every mess, however familiar it may seem, is in some way *sui generis*. And what this means is that our ability to apply our formalized knowledge to resolving the mess is always limited, because our know-how—that is, our ability to apply knowledge to practice with reliable consequences—is derived from an old context that is different from the new context. This is why another familiar aphorism—that generals are always fighting the last war—has so much bite. There’s a wonderful book by the historian Marc Bloch, called *Strange Defeat*, about how the French managed to get overrun by the Germans so quickly and easily in 1939. Bloch recounts how the French military leadership was simply unable to understand the implications of the new technologies of the German Blitzkrieg, technologies that literally changed the meaning of time in warfare. Operating under assumptions about time from the last war—about how long, for example, it would take a column of tanks to advance 20 miles—it was impossible for the French to understand what was happening to them until it was too late. The know-how so hard-won during the previous world war was useless.

Now a key point here is that human know-how as applied to complex, context-sensitive problems is not really cumulative—it does not improve very much with time. Rather, real-world practitioners are continually readjusting to deal with context-dependent complexity. One would be hard put to show that people today who are in positions where they must constantly deal with context-sensitive complexity, be they corporate CEOs, diplomats, politicians, generals, natural resource managers, or social activists, are systematically and systemically better at what they do than people in similar positions 50 or 100 years ago or even more. This despite all the important insights generated by social and natural science over that interval. We might get better at acting within a particular type of context, but the context is never fixed for very long. That’s why we must settle for muddling through.

But against this backdrop of complexity and messiness and muddling through and stationary know-how there is obviously something else going on. We actually do manage to improve our capacity to address some types of problems, even complex ones. The litany is perhaps boringly familiar. Public health, for example, is much, much better than it was a century ago. Modes of transport are inestimably more reliable and rapid. Ditto for modes of communication. Global agriculture provides more than enough nutrition for the world’s six-billion-plus people, an achievement that was literally unthinkable in the time of, say, Thomas Malthus—or even the Club of Rome. Obviously one of the main features of this type of improvement is that it usually reflects, and is an outgrowth of, technological change, but a related, less-remarked-upon feature is that it also reflects improved know-how, by which I mean improved ability to apply purposive action successfully to the achievement of particular goals.

So I want to emphasize this contrast, where our know-how about dealing with complex messes is, almost by definition, something that does not advance in a cumulative and decisive way; but our know-how in other domains, often those with a strong technological core, does advance significantly and sometimes rapidly as well, enabling us to solve particular problems that in the past seemed out of our reach.

Now a strong thread of analysis and critique among those who are interested in the relations among science, technology, and society has focused on the way that this technological capability creates new complexity of its own, and new types of messes that humans have no choice but to adapt to. As one obvious example, automobiles and power plants help get us to climate change, not to mention suburban sprawl. The rapid build-up of nuclear weapons in the Cold War, the mounting recognition, starting in the 1960s, of pervasive technological despoliation of the environment, the often disastrous consequences of technological “assistance” to developing nations, and the highly asymmetric distribution of technological power in the world, cumulate as an indictment of mindless techno-optimism.

This perspective tells us, quite correctly, that no technology is an island, that everything is connected to everything else, and that naïve efforts to solve complex problems through simple technological interventions will often bite back. We tar such interventions with the epithet “technological fix,” a term that indicates simplistic thinking about system complexity, and further suggests that technologies are being used as a short-cut around more meaningful and satisfying change in human behavior. I think a lot of the language around climate change has had this sort of flavor.

Yet this systems-level view may conceal something crucial.. Yes, technological systems often present the same sorts of unmanageable messiness that characterizes other complex systems. But messiness at the systems level may co-exist with extraordinary control and reliability at the level of day-to-day practice. One way to frame this discontinuity is to portray the control and reliability at the “shop floor” level of practice as an illusion that feeds human ambition and hubris and allows us to blunder, time and time again, into the unintended, emergent consequences of our own ingenuity. This, I think, is more-or-less the perspective that much critique of the technological society has taken in recent decades, and that underlies the common suspicion about “fixes.” And there is much truth here, of course.

But I want to explore a different perspective, where the discontinuity between technological systems and technologies-in-use is not a divide between the real world of complex messes and a hallucinogenic world of vain, reductionist ambition, but a divide between two different co-existing worlds, one of complexity and indeterminacy, and another that at times displays a very high degree of effectiveness in linking human intent to the achievement of particular outcomes. This potential for effective action has important political implications that I want to focus on here.

Consider two very different policy goals: the first, progressively improving the literacy levels of the citizens of the United States; and the second, eradicating smallpox from the

world. Both goals demand the application of know-how in very diverse contexts. Yet one of these areas of endeavor is characterized by ongoing frustration and failure, the other by remarkable success. Improving literacy in the world's most affluent nation might seem a rather straightforward, if not simple, task. Yet persistent effort has yielded little if any appreciable improvement in average reading levels among the nation's high school students, despite the strong national consensus on the importance of making progress in this area, and the considerable effort directed at improving knowledge about, and the practice of, teaching reading. In marked contrast, as everyone knows, smallpox has indeed been eradicated from the globe, despite what might have reasonably been seen as insuperable challenges, ranging from the difficulties of developing the necessary international cooperation regimes during the Cold War, to the complexities of administering vaccines in regions with little medical infrastructure and, at times, strong cultural aversion toward western medicine itself.

The key difference here is that one goal is served by a highly reliable core of action—the smallpox vaccine—and the other goal has no analogously effective core. To be sure, the effort to eradicate smallpox required adaptive institutional arrangements to allow deployment in a variety of cultural contexts. Yet this adaptation was aimed at deployment of a novel capability, embodied in the vaccine, which then produced a similar result wherever it was administered. The teaching of reading, on the other hand, which seems easy to do, has no similar reliable core of action—what works on children in one context (e.g., rich suburban private schools) may or may not work in another (e.g., poor urban public schools). Know-how about preventing smallpox can improve cumulatively because of the effectiveness of the vaccine itself, which provides a reliable baseline for learning. Vaccines can systematically be made more reliable and effective, and institutional arrangements for delivering them are not particularly important so long as the delivery occurs. Teaching reading has no comparably effective core of action—rather, it has numerous competing theories and methods. As a consequence, know-how remains shackled to context, and successes are not easily transferrable. Of course advocates of particular theories and methods of teaching reading insist that their approach would work if everyone used it, but in the absence of a context-independent core of reliable action, no particular set of claims can win out.

So I want to highlight several attributes of the vaccine that I suggest are generally typical of what I will call “effective things.” (I’m using this term for the moment to make clear that I’m not just referring to hard technological artifacts, but also, potentially, to routines, heuristics, procedures, and so on.) First, the cause-and-effect relations connecting the action (administering a vaccination) to the desired outcome (immunity from smallpox) are largely contained in, or embodied by, the thing itself. Second, the effectiveness of the thing is largely independent of context. The vaccine generally does its job regardless of who is administering it (as long as they have suitable training), or who is receiving it. Third, the cause-effect relations are often readily apparent. They do not require much interpretation, so disagreements about effectiveness can be fairly readily settled. Fourth, performance of the thing tends to improve incrementally over time, with continued use and experience.

What these attributes add up to is that an effective thing confers a conspicuous, rather context-independent advantage over other paths to a particular outcome. Effective things accomplish a task or further a goal better than the alternatives. Effective things, therefore, ought to help resolve disputes about alternative pathways to a particular goal. For example, if a computer program—or a brain implant—came along that with great reliability allowed the achievement of high levels of literacy in children regardless of context of use, we ought to expect that the contending parties in the battles over reading pedagogy would converge on that clearly effective approach. To not converge on that approach would be to choose ineffectiveness relative to this goal. People do, in fact, sometimes choose ineffectiveness. But in the case of vaccines, for example, while new vaccines are often met with opposition—for example, mandatory smallpox vaccination programs were greeted with riots in some U.S. cities at the turn of the 20th century—effective vaccines generally manage to attract a diverse constituency while opposition generally atrophies.

In fact, vaccination programs in the U.S. reflect an impressive degree of policy and operational coordination among diverse actors and institutions that in other contexts are typically at one-another's throats, including doctors, parents, school districts, insurance companies, vaccine manufacturers, and government regulatory bodies. This diverse group of actors and interests is every bit as complexly pluralistic as the group that continues to battle over competing approaches to teaching reading. The key difference in these two cases, I am arguing, is that one involves an effective technological core of action and the other does not.

Now a very important related point here is that one can scarcely imagine a bigger mess, one that implicates more competing interest groups and more operational complexity, than the medical system in the U.S.. Yet this mess co-exists with the remarkable shop-floor effectiveness of the vaccine itself. So we can have these two realities operating simultaneously, the reliability and effectiveness of vaccines, and the unmanageable complexity of the medical system as a whole.

So what I want to suggest is that increasing know-how and effectiveness *create the possibility* of political action on difficult and contentious problems. I want to consider that “political will” to solve such problems is catalyzed by the availability of know-how relevant to solving the problem, not vice versa. This, of course, is backwards to the way we typically think about such things, where the political will to act is usually advanced as a precondition to effective action. In this standard formulation, “political will” is created by shared awareness of the scope of the problem, awareness that is often enhanced and legitimated by scientific research aimed at documenting what is actually going on. Yet this common formula is bizarrely silent on how “political will” is supposed to translate into action in the absence of effective things to deploy that can help tame some aspect of the messiness of the problem.

Let me briefly illustrate the dilemma with an example from another domain. The story I want to mention tells how in the 1980s the nations of the world came to an agreement—the Montreal Protocol on Substances that Deplete the Ozone Layer—to phase out the

production of chlorofluorocarbons, a class of technologically and economically important refrigerants and solvents that also happen to destroy the stratospheric ozone layer that protects Earth from harmful ultraviolet radiation.

The popular narrative goes something like this: The CFC-ozone problem was discovered by basic scientific research on atmospheric chemistry; the results raised public and political concern; science evolved over the next 20 years and eventually stimulated an international response, the Montreal Protocol, which was first signed by 24 countries plus the EC in 1987 amidst continuing scientific controversy and opposition from many other countries. Conclusive scientific demonstration of the causal relations between CFC emissions and loss of stratospheric ozone over Antarctica came shortly afterwards, at which point opposition to the treaty from many nations and from the chemical industry disappeared. The treaty came into force in 1989 and by the early 1990s most nations of the world had signed on. The lesson?: Faced with definitive knowledge of a clear, shared risk, the nations took effective action. This is a story of science forcing right behavior, a story that provided the model upon which the response to climate change was later based.

But isn't there something missing here? Did the many nations that agreed to phase out CFC production decide to live without the benefits of keeping refrigerators and buildings cold, or keeping semiconductors clean? Of course not. The missing element here is the technological. In reports and articles about the Montreal Protocol, much is made about how confirmation of the link between CFCs and the Antarctic ozone hole led DuPont to immediately declare that it would stop producing CFCs, and made it impossible for the chemical industry as a whole to oppose the treaty. This version of events satisfies the perspective of both the scientific world, because we see facts stimulating rational action, and the commercial world, because we see corporations having no choice but to act responsibly in the face of evidence. But the chemical industry had been exploring CFC substitutes since the mid 1970s. By the mid-1980s, DuPont in particular came to realize not just that an array of alternatives were feasible, but that they offered a route to significant new profitability and competitive advantage.

My intent, by the way, is not to criticize DuPont; on the contrary, they were doing what firms are supposed to do. The real point is that the success of the Montreal Protocol was made possible because effective alternatives to CFCs were coming on-line. Science and diplomacy helped create the conditions for following this path. But had technological substitutes not existed, concerted global action would simply not have been possible, because society was highly dependent on the functions served by CFCs. There would have been plenty of room both for continued political conflict and continued scientific debate. Technological alternatives eliminated the need for disagreement.

The availability of alternatives to CFCs made it possible to meet the goals of multiple constituencies with conflicting values and worldviews: for example, 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.

So I'm proposing here a theory that links human rationality, technical capacity, and political effectiveness. I start with the presumption that when people make a decision, their intent is to achieve the objective at which their decision is aimed. At the institutional level, I suggest similarly that the adoption of a policy is intended to achieve the objectives at which the policy is aimed—although different people may of course interpret those objectives differently. Now if these rather modest claims to rationality have any grip, then there should be convergence in society toward “things that work,” that is, things that allow people to make decisions, and institutions to adopt policies, that are more successful at achieving the goals that they are meant to achieve than by using other approaches. And I've used the examples of vaccines and CFC alternatives to show how groups with very diverse interests and worldviews can indeed converge around something that does act more effectively in the world than other things.

Let me flesh this idea out further by applying it to what seems to be, for the moment at least, the great challenge of our generation, and I'm speaking of course about climate change. Now I understand that among many people concerned about climate change, the major obstacle to making progress is seen as the political intransigence of the U.S. But I think this is incorrect. It is of course true that the U.S. has been intransigent, but it is equally true that Europe, and the world, have made no significant progress in dealing with the climate problem, and I believe this in part reflects confusion about where, and how, we ought to expect progress to occur.

Now the mental model surrounding the formal approach to climate change follows the lessons of the Montreal Protocol in treating scientific research as the engine of necessary political change. Above all we must develop a comprehensive understanding of the fundamental behavior of the coupled ocean-atmosphere system, including the impacts of human activity on that system, and prediction of future evolution of the system, as reported every several years by the Intergovernmental Panel on Climate Change. This knowledge is what will compel action, which in this case primarily has meant action within the strictures of the UN Framework Convention on Climate Change and the Kyoto Protocol. Science proves that there is a problem, and demonstrates the need for action. This mental model has motivated the expenditure of enormous sums on climate science research, and on setting up an international governance regime aimed at mobilizing the nations of the world to cooperate in reducing emissions of gases that contribute to greenhouse warming.

Keep in mind that action—most importantly, the establishment of an appropriate market price for carbon emissions—is to be motivated by prediction of the bad things that will happen in the future as a result of our actions today—bad things like more floods and bigger hurricanes, longer droughts, stressed ecosystems and agricultural systems, and resurgent infectious diseases—amplification, that is, of complex messes that we already

face, and that we have confronted with highly variable attention and success absent the motivating fear of climate change. So the chain of logic here is that science will motivate nations to cooperatively establish policies that will force them to reduce their emissions by changing their behavior in a variety of ways, for example by using less energy and investing in energy technology research. The result of these actions will be a reduction of the future magnification of bad things.

Problems like natural disasters, biodiversity loss, declining availability of clean water, and resurgent infectious diseases are already very serious challenges to human well-being, challenges that have been growing worse by the decade, for reasons mostly of human development patterns such as urban and coastal population growth. Climate change will make many of the problems even worse, however. Scientific research on climate will motivate people and nations to take actions aimed at slowing this “even worse” part of the problem, reducing its effects progressively into the distant future. In other words, climate science will make more compelling the reasons for addressing a range of climate-related issues that already greatly challenge society, and which we have yet to address effectively. But, to be clear, when thinking about such challenges as natural disasters, biodiversity loss, or infectious diseases individually, the approaches to reducing vulnerability depend little, if at all, on scientific knowledge about climate change. Apparently, then, climate change knowledge points to the need to reduce emissions of greenhouse gases so that these problems won’t become even worse in the future, but it does not seem to advance action through other means.

So climate change discussions implicate, and confuse, two very different types of problems. The first is the problem of societal vulnerability to a variety of multi-causal challenges that we expect to have some growing connection to climate change—floods, droughts, and the like. The second is the problem of reducing the greenhouse gas emissions that will make these challenges even worse. Now these problems are distinct in two key dimensions. First, as a temporal matter, the inertia in both the global energy system and in the behavior of greenhouse gases in the atmosphere means that reduction of the magnifying effect of global warming on already existing problems cannot occur for many decades. Second, as matter of effective action in the world, reducing greenhouse gases is a different problem than reducing social vulnerability to climate impacts. And what has happened in the complex political world of climate change is that the moral challenge of reducing social vulnerability to climate impacts is being used to motivate action on the largely technological challenge of reducing carbon emissions. But, since these are two very different types of problems, the result is both confusion in thinking and confusion in action that lies at the heart of the failure of climate change policy to date.

In 1998, 5000 people died in Nicaragua in a matter of a few minutes from a mudflow triggered by Hurricane Mitch. Hurricanes are relatively common in Central America, but the problem was that the people were living on a deforested mountain slope. When the slope became saturated with water, it collapsed into a wave of mud.

The devastation of New Orleans by Hurricane Katrina in 2005 tells a similar tale. The

progressive development of the city and the environmental destruction of the surrounding wetlands rendered it increasingly vulnerable to hurricanes, while the levees that were designed to protect the city under precisely the circumstances that Katrina presented were poorly constructed and maintained. While there was enough suffering and loss to ensure that most everyone living around New Orleans got a good dose, it was the poor, the disenfranchised, the infirm, and the historically discriminated-against who suffered most and were disproportionately left behind to fend for themselves. Indeed, New Orleans, with its stark juxtaposition of the affluent and the poor, provided a synoptic portrayal of the global challenge.

Though the devastating consequences of these events cannot be attributed to climate change, they nevertheless became powerful symbols of it. In fact, even the 2004 South Asian tsunami, which had no connection to climate phenomena whatsoever, was used by influential voices to illustrate what a climate-changed-world might look like. Sir David King, Britain's chief science adviser, told the BBC at the time, "What is happening in the Indian Ocean underlines the importance of the Earth's system to our ability to live safely. And what we are talking about in terms of climate change is something that is really driven by our own use of fossil fuels."

Now I want to emphasize two attributes of these types of disasters. First, the level of scientific understanding surrounding both causes and impacts was very high. For example, scientists have known—and warned for decades—that the location of New Orleans on a rapidly subsiding river delta in the heart of the hurricane belt made some version of Katrina entirely inevitable. Similarly, the conditions leading to the Nicaraguan mudflow had been well-modeled and accurately predicted by scientists at the U.S. Geological Survey before the tragedy occurred.

Second, the natural disaster problem cannot be linked to a coherent technological solution path. Many well-tested policies are available to help reduce vulnerability to natural disasters. These range from building codes that can keep structures from collapsing in a storm, to land use regulations that limit construction in flood-prone areas, to environmental laws that preserve natural features, such as wetlands and forested slopes, that act as buffers against disasters. Yet all such policies are complex to adopt and implement, typically pitting vested interests against one another and demanding reasonably functional enforcement at local levels. And cause-effect relations are highly attenuated in time, and difficult to demonstrate. Disaster vulnerability is a messy problem that does not seem easily amenable to cumulative learning that can nucleate around effective technologies. The clearest path to reduced climate vulnerability seems to be increased and better-distributed wealth, and in this sense the problem is actually a subset of the larger messy problem of addressing global poverty. Yet as Katrina shows, even national wealth is not a very good predictor of effectiveness here.

The greenhouse gas problem is different—at least potentially. While carbon dioxide emissions have been rising progressively since the industrial revolution, these increases have also been accompanied by a progressive decarbonization of primary energy sources, and by a progressive decarbonization of economic activities in industrialized countries.

To be sure, because overall energy use and economic activity continue to rise, emissions continue to rise as well, and at an alarming rate—but the history of energy technology tells us that we can continue to decarbonize the energy system, and the history of technological innovation more generally—especially in the past fifty years or so—tells us that we can significantly increase the rate of decarbonization through an appropriate portfolio of investments and regulations. This is the type of problem that technologically sophisticated countries have learned how to solve. In the case of decarbonization, however, we have barely tried (though the widespread adoption of nuclear power by France perhaps puts them farther along the path). For example, G-7 nations have been *disinvesting* in energy research and development for the past 25 years or so, perhaps by as much as 65 percent—an extraordinary trend that is only now beginning to reverse itself. And the Kyoto Protocol is amazingly unaggressive on the question of catalyzing technological innovation, choosing instead to focus on targets and timetables aimed mostly at changing consumption behavior.

But the learning curves and the historical trends are in the right direction. The challenge here is to accelerate the process. Is this best done primarily through a combination of scientific evidence and moral arguments intended to overcome competing interests and get people to modify their lifestyles in particular ways so that they use less energy, which would at best have a small effect on emissions? Or through the potential of mobilizing technological effectiveness to satisfy multiple, competing interests in many different settings *without* demanding significant changes in values, interests, and ways of understanding the world?

In other words, while much of the climate impacts problem, especially in the area of disaster vulnerability, will have to be confronted by a muddling through process where cumulative learning and effectiveness is unlikely to be a major factor, the emissions reduction problem is quite plausibly going to be something that we can confront with increasing effectiveness—should we choose to take seriously the problem-solving potential of technological innovation. As with the example of continually improving public health within a dysfunctional medical system, the idea here is not to try to centrally manage an ungovernable global energy “system,” but to stimulate technological effectiveness, and thus political convergence, along as many independent paths as possible through appropriate investments and policies. It’s also worth remembering that emissions reduction paths can potentially be autocatalyzing due to the economic edge that increased energy efficiency may create for those nations that innovate most successfully in this domain.

Am I making an indiscriminate case for technological fixes? No, I am making a discriminate case. The question is: how to discriminate? One way to approach the problem is to consider whether the desired cause-effect relations are embodied in the technology itself or in a larger system. For example, technologies that capture carbon at the production source do not depend for their effectiveness on knowledge of how the climate system behaves. The desired effect is entirely attributable to the operation of the technology, and can be demonstrated at the site of operation. In contrast, most proposed geo-engineering fixes, such as iron fertilization of oceans, or planetary albedo

modification, depend for their success on comprehensive understanding of cause-effect relations within the open earth system. Biofuel is similarly problematic, in that the cause-effect relations relevant to decarbonization are distributed in the carbon cycle, not located in the technology itself. Approaches with distributed cause-effect chains are more likely to generate surprising and undesired consequences that do not further the very specific goal of reduced emissions. The move to biodiesel in Europe, for example, seems to be promoting tropical deforestation due to increased demand for palm oil. Hybrid vehicles, which isolate the decarbonization in the electric motor itself, appear to offer a sounder “fix.”

I am also not making a case for focusing investments and efforts on technological fixes at the expense of more difficult problems. It seems quite apparent that we are hugely underinvested in both research and on-the-ground action that could reduce exposure to climate impacts. But expectations for significant progress easily assessed over time scales even on the order of a decade or more should be modest. And I want to reiterate that the political failure in our overall approach to climate change has been on both ends of the continuum: we are acting as if vulnerability to climate can be solved by reducing emissions, which is wrong, and we are acting as if reducing emissions is essentially a problem of behavioral change in response to factual information, which is also wrong. These wrongs are compounded because they make the very hard problem—reducing vulnerability—seem easier than it is, and they have made the fairly straightforward (I won’t say easy) problem—reducing emissions—harder than it needs to be.

My own reflections on these questions, and the discomfort that they have created for me, can be captured by what one might call the Progressive’s Dilemma. If we perceive that a problem is caused by moral or ethical failure, then we want to solve that problem by correcting the failure, which means through behavioral change motivated by a clear understanding of both the moral and the factual elements of the problem. Solving a problem by addressing the underlying social causes seems both normatively and rationally more satisfactory than solving it by introducing technologies that get us off the hook for our sins. Of course, such a position is a commitment to both moral and scientific absolutism, with all the irony that such a joint commitment entails. More practically, it is also a commitment to long, hard political work whose outcome may never feel very satisfactory.

Technological fixes do not offer a route to moral or political absolutism. Consider the potential for technologies that capture carbon either at the source of emissions, such as a coal-fired power plant, or directly from the air. Neither of these approaches would violate my suggested guideline of not depending for their effectiveness on system-wide knowledge. But if successful, these types of approaches could rob us of an opportunity to struggle against the aspects of fossil fuel consumption that we find obnoxious. If we find ourselves regretting that opportunity, then it tells us that our concern about climate change is not simply one of its projected impacts, but of its cultural origins. My strong sense is that debates around climate change have not encouraged a clear separation of ends and means by the various constituencies involved, and that conflicts about cultural values have been playing out in the guise of conflicts about climate science and policy.

Putting technology in the center of the discussion could have a very salutary effect in this regard, by creating more clarity about desired ends and preferred means.

I am told, by someone who understands the state of research on machine-neural interfaces, that scientists are perhaps 20 years from being able to directly intervene in the brain to enhance various higher-level cognitive capacities, for example, the capacity to read with facility. Whether or not this is the case, it provides a thought experiment for moving the teaching of reading from the realm of political and scientific debate to the realm of technological effectiveness. Were we to achieve the capacity to download reading ability into any brain, we might regret the lost opportunity that lack of progress in teaching gives us to fight for better salaries for teachers, better home environments for children, or greater equity in wealth distribution. But of course we haven't been making much progress in those battles in the U.S. in any case, and it's doubtful that having a technological fix for teaching children to read would do much damage to those good causes—it might even advance them.

But one thing is clear: Should such technological fixes appear, constituencies holding diverse interests, values, and ways of knowing would surely aggregate around them, because they would recognize a potent tool for acting in the world. Groups that chose to not take advantage of the reliable effectiveness of the technology—because it offended their values, or contradicted their understandings of how the world works—would be marginalized and disempowered. This might be perfectly acceptable as a general principle to some groups—consider, for example, the Amish, or other tightly knit, technology-avoiding groups—by dint of the strength of their value systems and cohesiveness of their communities. But as a general matter, taking a principled stance against the use of an effective technology to resolve something widely acknowledged to be a problem is an inherently disempowering action. You are ceding effectiveness to others, whose interests and ways of understanding the world are now bound up—and advanced—in the use of that technology.

There are plenty of reasons to distrust claims that certain technologies will solve complex problems that have resisted solution in the past. There are probably thousands of software packages that promise to revolutionize the teaching of reading skills, but they have failed to do so, and will continue to fail, because they cannot tame the contextual complexity of the education process. But technological capabilities can also impose a simplifying order on problems that are complex in the absence of such capabilities. This emergent effectiveness is likely to be found on the “shop-floor,” at the level of practice. And it is likely to be accompanied by new domains of complexity at the system level. This mix of ground-level effectiveness and systems-level complexity is unavoidable; it is an aspect of the human condition.

For the past forty years or so, progressive politics has been generally suspicious of technological fixes, and with good reason. But I want to end with a suggestion that some notion of pragmatic technological progressivism needs to be resurrected as a part of any hopeful agenda for enhancing justice, equality, freedom, and even mutual understanding in the world. I've pointed to four reasons why such a resurrection ought to be

encouraged. First, the core reliability embodied in technologies is sometimes ideally suited for making progress on problems that are intractable when approached as political, behavioral, or moral problems. Second, in a world of finite attention and resources, we need to be smart in the way we choose to approach problems. We have not been particularly smart in the political approach to climate change, for example. Third, effective technologies can act as political attractors, bringing together diverse and even conflicting constituencies who recognize a common interest in the outcomes that can be reliably achieved. Technology, that is, may be mobilized as a powerful tool for conflict resolution, for example if applied to disputes over limited natural resources. Finally—and flowing from the previous reasons—a decision to abjure a technology is a decision to abdicate the effectiveness and therefore the power to achieve one’s aims that the technology confers. This seems to me like a very poor principle upon which to exercise political action.

Humans are an innovating species. 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 include this fundamental, if uncomfortable, reality of the human condition.

Acknowledgements: This talk draws heavily on ongoing collaborations with Richard R. Nelson, Helen Ingram, Ann Schneider, Steve Rayner, and Roger Pielke, Jr.