

ESSAYS / 01

POLICY MAKING IN THE POST-TRUTH WORLD:

*ON THE LIMITS OF SCIENCE AND THE
RISE OF INAPPROPRIATE EXPERTISE*

STEVE RAYNER AND DANIEL SAREWITZ

Steve Rayner died a couple of months after he and I finished up a baggy first draft of this essay and circulated it to a few colleagues. The essay itself was to be the first of several that we had been discussing for years about science, technology, politics, and society. So I had a thick folder-full of notes that I could draw on while making revisions, to assure myself that the end product was one Steve would have fully approved of, even if it was not nearly as good as we could have achieved together.

Had Steve not died shortly before the onset of the COVID-19 pandemic, we would certainly have given it a central role in this revised version. But I never had the benefit of Steve's insights on the strange unfolding of this disaster, and so, except in a couple of places where the extrapolation seems too obvious to not mention, the virus does not appear in what follows. Nonetheless, I know that Steve would

have relished an obvious irony related to expertise and the societal response to COVID-19: some experts proclaimed a welcome reawakening of public respect for experts triggered by the pandemic, even as other experts were insisting that the course of the disease marked a decisive repudiation of the legitimacy of experts in modern societies. Which seems as good an entry point as any into our exploration of the troubled state of expertise in today's troubled world.

* * * * *

Writing of his days as a riverboat pilot in *Life on the Mississippi*, Mark Twain described how he mastered his craft: “The face of the water, in time, became a wonderful book — a book that was a dead language to the uneducated passenger, but which told its mind to me without reserve, delivering its most cherished secrets as clearly as if it uttered them with a voice.”

The “wonderful book” to which Twain refers, of course, can nowhere be written down. The riverboat pilot’s expertise derives not from formal education but from constant feedback from his surroundings, which allows him to continually hone and test his skill and knowledge, expanding its applicability to a broadening set of contexts and contingencies. “It was not a book to be read once and thrown aside, for it had a new story to tell every day,” Twain continued. “Throughout the long twelve hundred miles there was never a page that was void of interest, never one that you could leave unread without loss.”

Expertise, in this way, necessarily involves the ability to make causal inferences (drawn, say, from the pattern of ripples on the surface of a river) that guide understanding and action to achieve better outcomes than could be accomplished without such guidance. Such special knowledge allows the expert to reliably deliver a desired outcome that cannot be assured by the nonexpert.

Expertise of this sort may also require lengthy formal training in sophisticated technical areas. But the expertise of the surgeon, or the airline pilot, is never just a matter of book learning; it requires the development of tacit

knowledge, judgment, and skills that come only from long practical experience and the feedback that such experience delivers. Expert practitioners demonstrate their expertise by successfully performing tasks that society values and expects from them, reliably and with predictable results. They navigate the riverboat through turbulent narrows; they repair the damaged heart valve; they land the aircraft that has lost power to its engines.

Yet every day it seems we hear that neither politicians nor the public are paying sufficient heed to expertise. The claim has become a staple of scholarly assertion, media coverage, and political argument. Commentators raise alarm at our present “post-truth” condition, made possible by rampant science illiteracy among the public, the rise of populist politics in many nations, and the proliferation of unverifiable information via the Internet and social media, exacerbated by mischievous actors such as Russia and the alt-right. This condition is said to result in a Balkanization of allegedly authoritative sources of information that in turn underlies distrust of mainstream experts and reinforces growing political division.

And still, despite this apparent turn away from science and expertise, few doubt the pilot or the surgeon. Or, for that matter, the plumber or the electrician. Clearly, what is contested is not all science, all knowledge, and all expertise, but particular kinds of science and claims to expertise, applied to particular types of problems.

Does population-wide mammography improve women’s health? It’s a simple question, still bitterly argued despite 50 years of mounting evidence. Is nuclear energy necessary to decarbonize global energy systems? Will missile defense systems work? Does Round-up cause cancer? What’s the most healthful type of diet? Or the best way to teach reading or math? For all of these questions, the answer depends on which expert you ask. Should face masks be worn outdoors in public places during the pandemic? Despite its relevance to the COVID-19 outbreak, this question has been scientifically debated for at least a century.

If the purpose of expertise applied to these sorts of questions is to help resolve them so that actions that are widely seen as effective can be pursued, then it would seem that the experts are failing. Indeed, these sorts of controversies have both proliferated and become ever more contested. Apparently, the type of expertise being deployed in these debates is different from the expertise of the riverboat pilot in the wheelhouse, or the surgeon in the operating room.

Practitioners like river pilots and surgeons can be judged and held accountable based on the outcomes of their decisions. Such a straightforward line of performance assessment can rarely be applied to experts who would advise policy makers on the scientific and technical dimensions of complex policy and political problems. Advisory experts of this sort are not acting directly on the problems they advise about. Even if their advice is taken, feedbacks on performance are often not only slow, but also typically incomplete, inconclusive, and ambiguous. Such experts are challenged to deliver anything resembling what we expect — and usually get — from our pilots, surgeons, and plumbers: predictable, reliable, intended, obvious, and desired outcomes.

1.

Nobody worries whether laypeople trust astrophysicists who study the origins of stars or biologists who study anaerobic bacteria that cluster around deep-sea vents. The wrangling among scientists who are debating, say, the reasons dinosaurs went extinct or whether string theory tells us anything real about the structure of the universe can be acrimonious and protracted, but it bears little import for anyone's day-to-day life beyond that of the scientists conducting the relevant research. But the past half-century or so has seen a gradual and profound expansion of science carried out in the name of directly informing human decisions and resolving disputes related to an expanding range of problems for democratic societies involving technology, the economy, and the environment.

If it can be said that there is a crisis of science and expertise and that we have entered a post-truth era, it is with regard to these sorts of problems, and to the claims science and scientific experts would make upon how we live and how we are governed.

Writing about the limits of science for resolving political disagreements about issues such as the risks of nuclear energy, the physicist Alvin Weinberg argued in an influential 1972 article that the inherent uncertainties surrounding such complex and socially divisive problems lead to questions being asked of science that science simply cannot answer. He coined the term “trans-science” to describe scientific efforts to answer questions that actually transcend science.

Two decades later, the philosophers Silvio Funtowicz and Jerome Ravetz more fully elucidated the political difficulties raised by trans-science as those of “post-normal” science, in which decisions are urgent, uncertainties and stakes are high, and values are in dispute. Their term defined a “new type of science” aimed at addressing the “challenges of policy issues of risk and the environment.” (Funtowicz and Ravetz used the term “post-normal” to contrast with the day-to-day puzzle-solving business of mature sciences that Thomas Kuhn dubbed “normal science” in his famous 1962 book, *The Structure of Scientific Revolutions*.)

What Funtowicz and Ravetz stressed was the need to recognize that science carried out under such conditions could not — in theory or practice — be insulated from other social activities, especially politics.

Demands on science to resolve social disputes accelerated as the political landscape in the 1960s and 70s began to shift from a primary focus on the opposition between capital and labor toward one that pitted industrial society against the need to protect human health and the environment, a shift that intensified with the collapse of the Soviet Union in the 1980s. Public concerns about air and water pollution, nuclear energy, low levels of chemical contamination and pesticide residues, food additives, and genetically

modified foods translated into public debates among experts about the magnitude of the problems and the type of policy responses, if any, that were needed. It is thus no coincidence that the 1980s and 90s saw “risk” emerge as the explicit field of competing claims of rationality.

As with the previous era of conflict between capital and labor, these disputes often mapped onto political divisions, with industrial interests typically aligning with conservative politics to assert low levels of risk and excessive costs of action, and interests advocating environmental protection aligning with regimes for which the proper role of government included regulation of industry to reduce risks, even uncertain ones, to public health and well-being.

As such conflicts proliferated, it was not much of a step to think that the well-earned authority of science to establish cause-effect relations about the physical and biological world might be applied to resolve these new political disputes. In much the same logical process that leads us to rely on the expertise of the riverboat pilot and cardiac surgeon, scientists with relevant expertise have been called upon to guide policy makers in devising optimal policies for managing complex problems and risks at the intersection of technology, the environment, health, and the economy.

But this logic has not borne out. Instead, starting in the 1970s, there has been a rapid expansion in health and environmental disputes, not-in-my-backyard protests, and concerns about environmental justice, invariably accompanied by dueling experts, usually backed by competing scientific assessments of potential or actual damage to individuals and communities. These types of disputes constitute an important dimension of today’s divisive national politics.

2.

Why has scientific expertise failed to meet the dual expectations created by the rise of scientific knowledge in the modern age and the impressive performance record of experts acting in other domains of technological society? The difficulties begin with nature itself.

The distinguished anthropologist Mary Douglas was wont to observe that nature is the trump card that can be played to win an argument even when time, God, and money have failed. The resort to nature as ultimate arbiter of disagreement is a central characteristic of the modern Western world. The debate between Burke and Paine in the 18th century over the origins of democratic legitimacy drew its energy from fundamentally conflicting claims about nature. A century later, J. S. Mill observed that “the word unnatural has not ceased to be one of the most vituperative epithets in the language.” This remains the case today. When we assert that something is only natural, we draw a line in the sand. We declare that it is simply the way things are and that no further argumentation can change that.

How does nature derive its voice in the political realm? In the modern world, nature speaks through science. Most people do not apprehend nature directly; they apprehend it via those experts who can speak and translate its language. Translated to the political realm, scientists who would advise policy making draw their legitimacy principally from the claim that they speak for nature. That expertise is ostensibly wielded to help policy makers distinguish that which is correct about the world from that which is incorrect, causal claims that are true from those that are false, and ultimately, policies that are better from those that are worse.

Yet when it comes to the complicated interface of technology, environment, human development, and the economy, political combatants have their own sciences and experts advocating on behalf of their own scientifically mediated version of nature. What is produced under such circumstances, Herbert Simon observed in 1983, is not ever more reliable knowledge, but rather “experts for the affirmative and experts for the negative.” Under these all-too-familiar conditions, science clearly must be doing something other than simply reporting upon well-established cause-and-effect inferences observed in nature. What, then, is it doing?

A key insight was provided in the work of ecologist C. S. Holling, who revealed the breadth and variety of scientists’ assumptions about how nature

works by describing the seemingly contradictory ecological management strategies adopted by foresters to address problems such as insect infestations or wildfires.

If foresters were conventionally rational, they would all do the same thing when given access to the same relevant scientific information. However, in the diverse forest management approaches that were actually implemented, Holling and colleagues detected “differences among the worldviews or myths of nature that people hold,” leading in turn to different scientific “explanations of how nature works and the [different] implication of those assumptions on subsequent policies and actions.”

One view of nature understands the environment to be favorable toward humankind. In this world, a benign nature re-establishes its natural order regardless of what humans do to their environment. This version of rationality encourages institutions and individuals to take a trial-and-error approach in the face of uncertainty. It is a view that requires strong proof of significant environmental damage to justify intervention that restricts economic development.

From another perspective, nature is in a precarious and delicate balance. Even a slight perturbation can result in an irreversible change in the state of the system. This view encourages institutions to take a precautionary approach to managing an ephemeral nature. The burden of proof, in this worldview, rests with those who would act upon nature.

A third view of nature centers around the uncertainties regarding causes and effects themselves. From this perspective, uncertainty is inherent, and the objective of scientific management is not to avoid any perturbation but to limit disorder via indicators, audits, and the construction of elaborate technical assessments to ensure that no perturbation is too great.

The point is not that any of these perspectives is entirely right or entirely wrong. Social scientists Schwarz and Thompson noted that “each of these

views of nature appears irrational from the perspective of any other,” reflecting what they term “contradictory certainties.” There can be no single, unified view of nature that can be expressed through a coherent body of science. In the post-normal context, when science is applied to policy making and decisions with potentially momentous consequences, scientists and decision-makers are always interpreting observations and data through a variety of pre-existing worldviews and frameworks that create coherence and meaning. Different myths of nature thus become associated with different institutional biases toward action.

Consider claims that we are collectively on the brink of overstepping “planetary boundaries” that will render civilization unsustainable. In the scientific journal *Nature*, Johan Rockström and his colleagues at the Stockholm Resilience Centre argue that “human actions have become the main driver of global environmental change,” that “could see human activities push the Earth system outside the stable environment state of the Holocene, with consequences that are detrimental or even catastrophic for large parts of the world.”

A review by Nordhaus et al. contests these claims, challenging the idea that these planetary boundaries constitute “non-negotiable” thresholds, interpreting them instead as rather arbitrary claims that for the most part don’t even operate at planetary scale. Similarly, Brook et al. conclude that “ecological science is unable to define any single point along most planetary continua where the effects of global change will cause abrupt shifts or transitions to qualitatively different regimes across the whole planet.” Strunz et al. argue that civilizational “collapse” narratives are themselves subject to interpretation and that the supposed alternatives of “sustainability or collapse” mischaracterize not only the nature of environmental challenges, but the types of policy responses available to societies.

These various expert perspectives beautifully display the competing rationalities mapped out by Holling a generation before. They suggest that rather than non-negotiables, humanity faces a system of trade-offs — not

only economic, but moral and aesthetic as well. Deciding how to balance these trade-offs is a matter of political contestation, not scientific fact. What counts as “unacceptable environmental change” involves judgments concerning the value of the things to be affected by the potential changes.

Seldom do scientists or laypeople consciously reflect on the underlying assumptions about the nature of nature that inform their arguments. Even when such assumptions can be made explicit, as Holling discussed in the case of forest ecosystem management, it is not possible to say which provides the best foundation for policy making. This is the case given both that the science is concerned with the future states of open, complex, non-determinate natural and social systems, and that people may reasonably disagree about the details of a desirable future as well as the best pathways of getting there.

Amidst such multi-level uncertainty and disagreement (which may last for decades, or forever), it is impossible to test causal inferences at large enough temporal and spatial scales to draw conclusions about which experts were right and which were wrong with regard to questions related to something like overall earth-system behavior. Experts participating in such debates thus need never worry that they will be held accountable for the consequences of acting on their advice. They wield their expertise with impunity.

3.

The most powerful ammunition that experts can deploy are numbers. Indeed, we might say that if nature is a political trump card, numbers are what give that card its status and authority. Pretty much any accounting of science will put quantification at the center of science’s power to portray the phenomena that it seeks to understand and explain. As Lord Kelvin said in 1883: “When you can measure what you are speaking about and express it in numbers, you know something about it: but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind.”

To describe something with a number is to make a sharp claim about the correspondence between the idea being conveyed and the material world out there. Galileo said that mathematics was the language of the universe. The use of numbers to make arguments on behalf of nature thus amounts to an assertion of superior — that is, expert — authority over other ways to make claims about knowledge, truth, and reality.

When we look at the kinds of numbers that often play a role in publicly contested science, however, we see something surprising. Many numbers that appear to be important for informing policy discussions and political debates describe made-up things, not actual things in nature. They are, to be sure, abstractions about, surrogates for, or simulations of what scientists believe is happening — or will happen — in nature. But they are numbers whose correspondence to something real in nature cannot be tested, even in theory.

Yet even when representing abstractions or poorly understood phenomena, numbers still appear to communicate the superior sort of knowledge that Lord Kelvin claimed for them, giving rise to what mathematician and philosopher Alfred North Whitehead in 1929 termed “the Fallacy of Misplaced Concreteness,” in which abstractions are taken as concrete facts. This fallacy is particularly seductive in the political context, when complicated matters (for example, the costs versus benefits of a particular policy decision) can be condensed into easily communicated numbers that justify a particular policy decision, such as whether or not to build a dam or protect an ecological site.

Consider efforts to quantify the risks of high-level nuclear waste disposal in the United States and other countries. The behavior of buried nuclear waste is determined in part by the amount of water that can reach the disposal site and thus eventually corrode the containment vessel and transport radioactive isotopes into the surrounding environment.

One way to characterize groundwater flow is by calculating a variable called

percolation flux, an expression of how fast a unit of water flows through rocks, expressed in distance per unit of time. The techniques used to assign numbers to percolation flux depend on hydrogeological models, which are always incomplete representations of real conditions, and laboratory tests on small rock samples, which cannot well represent the actual heterogeneity of the disposal site. Based on these calculations, assessments of site behavior then adopt a value of percolation flux for the entire site for purposes of evaluating long-term risk.

Problems arise, though, because water will behave differently in different places and times depending on conditions (such as varying density of fractures in the rocks, or connectedness between pores, or temperature). At Yucca Mountain, Nevada, the site chosen by Congress in 1987 to serve as the US high-level civilian nuclear waste repository, estimates of percolation flux made over a period of 30 years have varied from as low as 0.1 mm/yr to as much as 40 mm/yr. This three-orders-of-magnitude difference has momentous implications for site behavior, with the low end helping to assure decision-makers that the site will remain dry for thousands of years and the high end indicating a level of water flow that, depending on site design, could introduce considerable risk of environmental contamination over shorter time periods.

To reduce uncertainties about the hydrogeology at Yucca Mountain, scientists proposed to test rocks from near the planned burial site, 300 meters underground, for chlorine 36 (^{36}Cl). This radioisotope is rare in nature but is created during nuclear blasts and exists in higher abundance in areas where nuclear weapons have been tested, such as the Nevada Test Site near Yucca Mountain. If excess ^{36}Cl could be found at the depth of the planned repository, it would mean that water had travelled from the surface to the repository depth in the 60 or so years since weapons tests were conducted, requiring a much higher percolation flux estimate than if no ^{36}Cl was present.

But confirming the presence of excess ^{36}Cl hinged on the ability to detect it at concentrations of parts per 10 billion, a level of precision that turned

out to introduce even more uncertainties to the percolation flux calculation. Indeed, contradictory results from scientists working in three different laboratories made it impossible to confirm whether or not the isotope was present in the sampled rocks. This research, performed to reduce uncertainty, actually increased it, so that the range of possible percolation flux values fully encompassed the question of whether or not the site was “wet” or “dry,” and fully permitted positions either in support of or opposing the burial of nuclear waste at Yucca Mountain.

And even if scientists were to agree on some “correct” value of percolation flux for the repository site, it is only one variable among innumerable others that will influence site behavior over the thousands of years during which the site must safely operate. Percolation flux thus turns out not to be a number that tells us something true about nature. Rather, it is an abstraction that allows scientists to demonstrate their expertise and relevance and allows policy makers to claim that they are making decisions based upon the best available science, even if that science is contradictory and can justify any decision.

Such numbers proliferate in post-normal science and include, for example, many types of cost-benefit ratio; rates and percentages of species extinction; population-wide mortality reduction from dietary changes; ecosystem services valuation and, as we will later discuss, volumes of hydrocarbon reserves. Consider a number called “climate sensitivity.” As with percolation flux, the number itself — often (but not always) defined as the average atmospheric temperature increase that would occur with a doubling of atmospheric carbon dioxide — corresponds to nothing real, in the sense that the idea of a global average temperature is itself a numerical abstraction that collapses a great diversity of temperature conditions (across oceans, continents, and all four seasons) into a single value.

The number has no knowable relationship to “reality” because it is an abstraction and one applied to the future no less — the very opposite of what Lord Kelvin had in mind in extolling the importance of quantification. Yet

it has come to represent a scientific proxy for the seriousness of climate change, with lower sensitivity meaning less serious or more manageable consequences and higher values signaling a greater potential for catastrophic effects. Narrowing the uncertainty range around climate sensitivity has thus been viewed by scientists as crucially important for informing climate change policies.

Weirdly, though, the numerical representation of climate sensitivity remained constant over four decades — usually as a temperature range of 1.5°C to 4.5°C— even as the amount of science pertaining to the problem has expanded enormously. Starting out as a back-of-the-envelope representation of the range of results produced by different climate models from which no probabilistic inferences could be drawn, climate sensitivity gradually came to represent a probability range.

Most recently, for example, an article by Brown and Caldeira reported an equilibrium climate sensitivity value of 3.7°C with a 50 percent confidence range of 3.0°C to 4.2°C, while a study by Cox et al. reported a mean value of 2.8°C with a 66 percent confidence range of 2.2°C to 3.4°C, and an assessment by Sherwood and a team of 24 other scientists reported a 66 percent probability range of 2.6°C to 3.9°C. The 2020 Sherwood team study characterized the initial 1.5°C to 4.5°C range, first published in a 1979 National Research Council report, as “prescient” and “based on very limited information.” In that case, one might reasonably wonder about the significance of four decades of enormous subsequent scientific effort (the Sherwood paper cites more than 500 relevant studies) leading to, perhaps, a slightly more precise characterization of the probability range of a number that is an abstraction in the first place.

The legacy of research on climate sensitivity is thus remarkably similar to that of percolation flux: decades of research and increasingly sophisticated science dedicated to better characterizing a numerical abstraction that does not actually describe observable phenomena, with little or no change in uncertainty.

4.

Ultimately, most political and policy disputes involve the future — what it should look like and how to achieve that desired state. Scientific expertise is thus often called upon to extrapolate from current conditions to future ones. To do so, researchers often construct numerical representations of relevant phenomena that can be used to extrapolate from current conditions to future ones. These representations are called models.

Pretty much everyone is familiar with how numerical models can be used to inform decision-making through everyday experience with weather forecasting. Weather forecasting models are able to produce accurate forecasts up to about a week in advance. In part, this accuracy can be achieved because, for the short periods involved, weather systems can be treated as relatively closed, and the results of predictions can be evaluated rigorously. Weather forecasts have gotten progressively more accurate over decades because forecasters make millions of forecasts each year that they test against reality, allowing improved model performance due to continual learning from successes and mistakes and precise measurement of predictive skill.

But that's not all. A sophisticated and diverse enterprise has developed to communicate weather predictions and uncertainties for a variety of users. Organizations that communicate weather information understand both the strengths and weaknesses of the predictions, as well as the needs of those who depend on the information. Examples range from NOAA's Digital Marine Weather Dissemination System for maritime users, to the online hourly forecasts at Weather.com.

Meanwhile, people and institutions have innumerable opportunities to apply what they learn from such sources directly to decisions and to see the outcomes of their decisions — in contexts ranging from planning a picnic to scheduling airline traffic. Because people and institutions are continually acting on the basis of weather forecasts, they develop tacit knowledge that allows them to interpret information, accommodate uncertainties, and

develop trust based on a shared understanding of benefits. Picnickers, airline executives, and fishers alike learn how far in advance they should trust forecasts of severe weather in order to make decisions whose stakes range from the relatively trivial to the truly consequential.

Even though the modeling outputs often remain inexact and fundamentally uncertain (consider the typical “50 percent chance of rain this afternoon” forecast) and specific short-term forecasts often turn out to be in error, people who might question the accuracy or utility of a given weather forecast are not accused of “weather science denial.” This is because the overall value of weather information is well integrated into the institutions that use the predictions to achieve desired benefits.

The attributes of successful weather forecasting are not, and cannot be, built into the kinds of environmental and economic models used to determine causal relations and predict future conditions in complex natural, technological, and social systems. Such models construct parallel alternative worlds whose correspondence to the real world often cannot be tested. The models drift away from the direct connection between science and nature, while giving meaning to quantified abstractions like percolation flux and climate sensitivity, which exist to meet the needs of modeled worlds but not the real one.

For example, the Chesapeake Bay Program (CBP), established in 1983, launched an extensive ecosystem modelling program to support its goal of undoing the negative effects of excessive nutrient loading in the bay from industrial activity, agricultural runoff, and economic development near the shoreline. A distributed suite of linked models was developed so scientists could predict the ecosystem impact of potential management actions, including improving sewage treatment, controlling urban sprawl, and reducing fertilizer or manure application on agricultural lands.

While the CBP model includes data acquired from direct measurements in nature, the model itself is an imaginary world that stands between science and nature. The difference between a modelled prediction of, say, decreased

nitrogen load in the Chesapeake Bay and an observation of such a decrease is that the achievement of that outcome in the model occurred by tweaking model inputs, parameters, and algorithms, whereas in nature the outcome was created by human decisions and actions.

And indeed, based on models that simulated the results of policy interventions, CPB claimed that it had achieved a steady improvement in the water quality of the main stem of the Bay. Yet interviews conducted in 1999 with program employees revealed that actual field testing did not demonstrate a trend of improved water quality. The computer model, designed to *inform* management of a real-world phenomenon, in fact became the *object* of management.

A similar phenomenon of displacing reality with a simulation occurs in modelling for climate policy when the impacts of nonexistent technologies — such as bioenergy with carbon capture and storage or solar radiation management interventions — are quantified and introduced into models as if they existed. Their introduction allows models to be tweaked to simulate reductions in future greenhouse warming, which are then supposed to become targets for policy making.

As with the Chesapeake Bay model, these integrated assessment models depend on hybrid and constructed numbers to generate concrete predictions. To do so, they must assume future atmospheric composition, land cover, sea surface temperature, insolation, and albedo, not to mention the future of economic change, demographics, energy use, agriculture, and technological innovation. Many of the inputs themselves are derived from still other types of models, which are in turn based on still other sets of assumptions.

Based on these models, some scientists claim that solar radiation management techniques will contribute to global equity; others claim the opposite. In fact, the models upon which both sets of claims depend provide no verifiable knowledge about the actual world and ignore all of the scientific, engineering, economic, institutional, and social complexities that will determine

Mythical Beasts of Policy-making



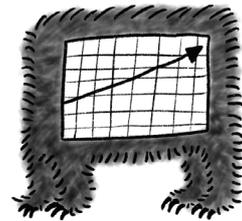
The Perfect
Externality Price



Unified and Exuberant
Stakeholders



The Utility-maximizing
Consumer



The Business-as-usual
Scenario

Samuel

real outcomes associated with whatever it is that human societies choose to do or not do.

The contrast between weather and climate forecasting could not be clearer. Weather forecasts are both reliable and useful because they predict outcomes in relatively closed systems for short periods with immediate feedback that can be rapidly incorporated to improve future forecasts, even as users (picnickers, ship captains) have innumerable opportunities to gain direct experience with the strengths and limits of the forecasts.

Using mathematical models to predict the future global climate over the course of a century of rapid sociotechnical change is quite another matter. While the effects of different development pathways on future atmospheric greenhouse gas concentrations can be modeled using scenarios, there

is no basis beyond conjecture for assigning relative probabilities to these alternative futures. There are also no mechanisms for improving conjectured probabilities because the time frames are too long to provide necessary feedback for learning. What's being forecast exists only in an artificial world, constituted by numbers that correspond not to direct observations and measurements of phenomena in nature, but to an assumption-laden numerical representation of that artificial world.

The problem is not by any means limited to climate models. Anyone who has followed how differing interpretations of epidemiological models have been used to justify radically different policy choices for responding to the COVID-19 pandemic will recognize the challenges of extrapolating from assumption-laden models to real-world outcomes. Similar difficulties have been documented in policy problems related to shoreline engineering, mine waste cleanup, water and fisheries management, toxic chemical policy, nuclear waste storage (as discussed), land use decisions, and many others.

And yet, because such models are built and used by scientists for research that is still called science and produce crisp numbers about the artificial worlds they simulate, they are often subject to Whitehead's fallacy of misplaced concreteness and treated as if they represent real futures. Their results are used by scientists and other political actors to make claims about how the world works and, therefore, what should be done to intervene in the world.

In this sense, the models serve a role similar to goat entrails and other pre-scientific tools of prophecy. They separate the prophecy itself, laden with inferences and values, from the prophet, who merely reports upon what is being foretold. The models become political tools, not scientific ones.

5.

When decisions are urgent, uncertainties and stakes are high, and values are in dispute — the post-normal conditions of Funtowicz and Ravetz — it turns out that science's claim to speak for nature, using the unique precision of

numbers and the future-predicting promise of models, is an infinitely contestable basis for expertise and its authority in the political realm.

And yet, science undoubtedly does offer an incomparably powerful foundation not only for understanding our world but also for reliably acting in it.

That foundation depends upon three interrelated conditions that allow us to authoritatively establish causal relationships that can guide understanding and effective action — conditions very different from those we have been describing, and with very different consequences in the world.

First is control: the creation or exploitation of closed systems, so that important phenomena and variables involved in the system can be isolated and studied. Second is fast learning: the availability of tight feedback loops, which allow mistakes to be identified and learning to occur because causal inferences can be repeatedly tested through observations and experiments in the controlled or well-specified conditions of a more or less closed system. Third is clear goals: the shared recognition or stipulation of sharply defined endpoints toward which scientific progress can be both defined and assessed, meaning that feedback and learning can occur relative to progress toward agreed-upon outcomes that confirm the validity of what is being learned.

Technology plays a dual role in the fulfillment of these three conditions. Inventions that observe or measure matter, such as scales, telescopes, microscopes, and mass spectrometers, translate inputs from nature into interpretable signals (measurements, images, waveforms, and so on) that allow scientists to observe and often quantify components and phenomena of nature that would otherwise be inaccessible. At the same time, the development and use of practical technologies such as steam engines, electric generators, airfoils (wings), cathode ray tubes, and semiconductors continually raise questions for scientists to explore about the natural phenomena that the technologies embody (e.g., the transformation of heated water into pressurized steam; the flow of fluids or electrons around or through various

media) and, in turn, derive generalizable relationships.

Under these three technologically mediated conditions, the practical consequences of scientific advances have helped to create the technological infrastructure of modernity. Technology, it turns out, is what makes science real for us. The light goes on, the jet flies, the child becomes immune. From such outcomes, people reliably infer that the scientific account of phenomena must be true and that the causal inferences derived from them must be correct. Otherwise, the technologies would not work.

Thus, our sense of science's reliability is significantly created by our experience with technology. Moreover, technological performance shares this essential characteristic with practitioner expertise: nonexperts can easily see whether this process of translation is actually taking place and doing what's expected. Indeed, expert practice typically involves the use of technology (a riverboat, a plumber's torch, a surgical laser) to achieve its goal.

The problem for efforts to apply scientific expertise to complex social problems is that the three conditions mostly do not pertain. The systems being studied – the climate-energy system, fluids in the earth's crust, population-wide human health – are open, complex, and unpredictable. Controlled experiments are often impossible, so feedback that can allow learning is typically slow, ambiguous, and debatable. Perhaps most importantly, endpoints often cannot be sharply defined in a way that allows progress to be clearly assessed; they are often related to identifying and reducing risk, and risk is an inherently subjective concept, always involving values and worldviews.

In the case of weather forecasts, vaccines, and surgical procedures, experts can assure us of how a given action will lead to a particular consequence, and, crucially, we can see for ourselves if they are right. In the case of science advisory expertise, the outcomes of any particular decision or policy are likely to be unknown and unknowable. No one can be held to account. Assumptions about the future can be modified and modeled to suit competing and conflicting interests, values, and beliefs about how the future ought

to unfold. Science advisory experts can thus make claims with political content while appearing only to be speaking the language of science.

The exercise of expert authority under such circumstances might be termed “inappropriate expertise.” Its origins are essentially epistemological: climate models, or the statistical characterization of a particular chemical’s cancer-causing potential, manifest a different type of knowledge than weather forecasts, jet aircraft, and vaccines. Claims to expertise based upon the former achieve legitimacy by borrowing the well-earned authority of the latter. In stark contrast, the legitimacy of expert-practitioners derives directly from proven performance in the real world.

6.

When we apply the authority of normal science to post-normal conditions, a *mélange* of science, expertise, and politics is the usual result. Neither more research nor more impassioned pleas to listen to and trust an undifferentiated “science” will improve the situation because it is precisely the proliferation of post-normal science and its confusion with normal science that are the cause of the problem.

Yet, in the face of controversies regarding risk, technology, and the environment, the usual remedy is to turn things over to expert organizations like the National Academy of Sciences or the UK Royal Society. But doing so typically obscures the normative questions that lie at the heart of conflicts in question. Why would anyone think that another 1,000 studies of climate sensitivity would change the mind of a conservative who opposes global governance regimes? Or that another decade of research on percolation flux might convince an opponent of nuclear power that nuclear waste can be safely stored for 10,000 years? Disagreements persist. More science is poured into the mix. Conflicts and controversies persist indefinitely.

There is an alternative. Decision-makers tasked with responding to controversial problems of risk and society would be better served to pursue

solutions through institutions that can tease out the legitimate conflicts over values and rationality that are implicated in the problems. They should focus on designing institutional approaches that make this cognitive pluralism explicit, and they should support activities to identify political and policy options that have a chance of attracting a diverse constituency of supporters.

Three examples from different domains at the intersection of science, technology, and policy can help illuminate this alternative way of proceeding. Consider first efforts to mitigate the public health consequences of toxic chemical use and exposure. Such efforts, in particular via the federal Toxics Substances Control Act (TSCA) of 1976, have historically attempted to insulate the scientific assessment of human health risks of exposures to chemicals from the policy decisions that would regulate their use. But from TSCA's inception, it has been clear that there is no obvious boundary that separates the science of risk from the politics of risk. The result — consistent with our discussion so far — has been endless legal action aimed at proving or disproving that scientific knowledge generated by EPA in support of TSCA was sufficient to allow regulation.

Starting in the late 1980s, the state of Massachusetts adopted an alternative approach. Rather than attempting to use scientific risk assessment to ban harmful chemicals that are valued for their functionality and economic benefit, Massachusetts' Toxics Use Reduction Act (TURA) of 1989 focused on finding replacements that perform the same functions. The aim was to satisfy the concerns of both those aiming to eliminate chemicals in the name of environmental health and those using them to produce economic and societal value.

TURA turned the standard adversarial process into a collaborative one. State researchers tested substitutes for effectiveness and developed cost-benefit estimates; they worked with firms to understand barriers to adoption and cooperated with state agencies and professional organizations to demonstrate the alternatives. Rather than fighting endless scientific and

regulatory battles, firms that use toxic chemicals became constituents for safer chemicals.

Between 1990 and 2005, Massachusetts firms subject to TURA requirements reduced toxic chemical use by 40 percent and their on-site releases by 91 percent. Massachusetts succeeded not by trying to reduce scientific uncertainty about the health consequences of toxic chemicals in an effort to compel regulatory compliance, but by searching for solutions that satisfied the beliefs and interests of competing rationalities about risk.

A second example draws from ongoing efforts to assess hydrocarbon reserves. In the 1970s and 1980s, coincident with national and global concerns about energy shortages, the US Geological Survey (USGS) began conducting regular assessments of the size of US hydrocarbon (oil and gas) reserves. As with percolation flux and climate sensitivity, quantified estimates of the volume of natural gas or oil stored in a particular area of the Earth's crust have no demonstrable correspondence to anything real in the world. The number cannot be directly measured, and it depends on other variables that change with time, such as the state of extraction technologies, the state of geological knowledge, the cost of various energy sources, and so on.

USGS assessments that reserves were declining over time were largely noncontroversial until 1988, when the natural gas industry began lobbying the government to deregulate natural gas markets. When the USGS assessment released that year predicted a continued sharp decline in natural gas reserves, the gas industry vociferously disagreed.

According to the American Gas Association (AGA), "The U.S. Geological Survey's erroneous and unduly pessimistic preliminary estimates of the amount of natural gas that remains to be discovered in the United States . . . is highly inaccurate and clearly incomplete . . . the direct result of questionable methodology and assumptions." In the standard ritual, dueling numbers were invoked, with the USGS report estimating recoverable natural gas reserves at 254 trillion cubic feet and the AGA at 400 trillion.

The customary prescription for resolving such disputes, of course, would be to do more research to better characterize the numbers. But in this case, the USGS adopted a different approach. It expanded the institutional diversity of the scientists involved in the resource assessment exercises, adding industry and academic experts to a procedure that had previously been conducted by government scientists alone.

The collective judgment of this more institutionally diverse group resulted in significantly changed assumptions about, definitions of, and criteria for estimating hydrocarbon reserves. By 1995, the official government estimate for US natural gas reserves went *up* more than fourfold, to 1,075 trillion cubic feet.

Agreement was created not by insulating the assessment process from stakeholders with various interests in the outcome, but by bringing them into the process and pursuing a more pluralistic approach to science. Importantly, the new assessment numbers could be said to be more scientifically sound only insofar as they were no longer contested. Their accuracy was still unknowable. But agreement on the numbers helped to create the institutional and technological contexts in which recovering significantly more oil and gas in the United States became economically feasible.

Finally, consider the role of complex macroeconomic models in national fiscal policy decisions. Economists differ on their value, with some arguing that they are essential to the formulation of monetary policy and others arguing that they are useless. Among the latter, the Nobel Prize-winning economist Joseph Stiglitz asserts: "The standard macroeconomic models have failed, by all the most important tests of scientific theory."

In the end, it doesn't appear to matter much. In the United States, the models are indeed used by the Federal Reserve to support policy making. Yet the results appear not to be a very important part of the system's decision processes, which depend instead on informed judgement about the state of the economy and debate among its governors. Whatever role the models

might play in the Federal Reserve decision process, it is entirely subservient to a deliberative process that amalgamates different sources of information and insight into narratives that help make sense of complex and uncertain phenomena.

Indeed, the result of the Federal Reserve's deliberative process is typically a decision either to do nothing or to tweak the rate at which the government loans money to banks up or down by a quarter of a percent. The incremental nature of the decision allows for feedback and learning, assessed against an endpoint mandated by Congress: maximum employment and price stability. The role of the models in this process seems mostly to be totemic. Managing the national economy is something that experts do, and using complicated numerical models is a symbol of that expertise, inspiring confidence like the stethoscope around a doctor's neck.

Each of these examples offers a corrective to the ways in which science advice typically worsens sociotechnical controversies.

The Federal Reserve crunches economic data through the most advanced models to test the implications of various policies for future economic performance. And then its members, representing different regions and perspectives, gather to argue about whether to take some very limited actions to intervene in a complex system — the national economy — whose behavior so often evades even short-term prediction.

When the US Geological Survey found itself in the middle of a firestorm of controversy around a synthetic number representing nothing observable in the natural world, it did not embark upon a decades-long research program to more precisely characterize the number. It instead invited scientists from other institutions, encompassing other values, interests, and worldviews, into the assessment process. This more cognitively diverse group of scientists agreed to new assumptions and definitions for assessing reserves and arrived at new numbers that would have seemed absurd to the earlier,

more homogeneous group of experts, but that met the information needs of a greater diversity of users and interests.

Toxic chemical regulation in the United States has foundered on the impossibility of providing evidence of harm sufficiently convincing to withstand legal opposition. More research and more experts have helped to enrich lawyers and expert witnesses, but they failed to restrict the use of chemicals that are plausibly but not incontrovertibly dangerous. The state of Massachusetts pursued a different approach, working within the uncertainties, to find complementarities between the interests and risk perspectives of environmentalists and industry in the search for safer alternatives to chemicals that were plausibly harmful.

Truth, it turns out, often comes with big error bars, and that allows space for managing cognitive pluralism to build institutional trust. The Federal Reserve maintains trust through transparency and an incremental, iterative approach to decision-making. The USGS restored trust by expanding the institutional and cognitive diversity of experts involved in its assessment process. Massachusetts created trust by taking seriously the competing interests and rationalities of constituencies traditionally at each other's throats.

Institutions are what people use to manage their understanding of the world and determine what information can be trusted and who is both honest and reliable. Appropriate expertise emerges from institutions that ground their legitimacy not on claims of expert privilege and the authority of an undifferentiated "science," but on institutional arrangements for managing the competing values, beliefs, worldviews, and facts arrayed around such incredibly complex problems as climate change or toxic chemical regulation or nuclear waste storage. Appropriate expertise is vested and manifested not in credentialed individuals, but in institutions that earn and maintain the trust of the polity. And the institutional strategies available for managing risk-related controversies of modern technological societies may be as diverse as the controversies themselves.

7.

We do not view it as coincidental that concerns among scientists, advocates, and others about post-truth, science denial, and so on have arisen amidst the expenditure of tens of billions of dollars over several decades by governments and philanthropic foundations to produce research on risk-related political and policy challenges. These resources, which in turn incentivized the creation of many thousands of experts through formal academic training in relevant fields, have created a powerful political constituency for a particular view of how society should understand and manage its technological, environmental, health, and other risks: with more science, conveyed to policy makers by science advocacy experts, to compel rational action.

Yet the experience of unrelenting and expanding political controversies around the risks of modernity is precisely the opposite outcome of what has been promised. Entangling the sciences in political disputes in which differing views of nature, society, and government are implicated has not resolved or narrowed those disputes, but has cast doubt upon the trustworthiness and reliability of the sciences and experts who presume to advise on these matters. People still listen to their dentists and auto mechanics. But many do not believe the scientists who tell them that nuclear power is safe, or that vaccines work, or that climate change is real.

We don't think that's a perverse or provocative view, but an empirically grounded perspective on why things haven't played out as promised. When risks and dilemmas of modern technological society become subject to political and policy action, doing more research to narrow uncertainties and turning to experts to characterize what's going on as the foundation for taking action might seem like the only rational way to go. But under post-normal conditions, in which decisions are urgent, uncertainties and stakes are high, and values are in dispute, science and expertise are, at best, only directly relevant to one of those four variables — uncertainty — and even there, the capacity for making a difference is often, as we've shown, modest at best.

The conditions for failure are thus established. Advocates and experts urgently proclaim that the science related to this or that controversy is sufficiently settled to allow a particular political or policy prescription — the one favored by certain advocates and experts — to be implemented. Left out of the formula are the high stakes and disputed values. Who loses out because of the prescribed actions? Whose views of how the world works or should work are neglected and offended?

Successfully navigating the divisive politics that arise at the intersections of technology, environment, health, and economy depends not on more and better science, nor louder exhortations to trust science, nor stronger condemnations of “science denial.” Instead, the focus must be on the design of institutional arrangements that bring the strengths and limits of our always uncertain knowledge of the world’s complexities into better alignment with the cognitive and political pluralism that is the foundation for democratic governance — and the life’s blood of any democratic society. //

Acknowledgments: *Heather Katz assured me that this is what Steve would have wanted me to do, and Jerry Ravetz assured me that this is what Steve would have wanted us to say. Ted Nordhaus helped me figure out how best to say it.*