I've been asked to talk
"for about 7 minutes" on

Risks Analysis of Emerging Technologies with High Uncertainty

and also about

Overconfidence

and to do it with 5 slides.

I'll stick to the 7-minute time constraint but violate the 5-slide constraint.

M. Granger Morgan (NAS Sec 64)
Department of Engineering and Public Policy
Carnegie Mellon University
https://engineering.cmu.edu/directory/bios/morgan-m-granger.html

Here is a standard framework for risk analysis

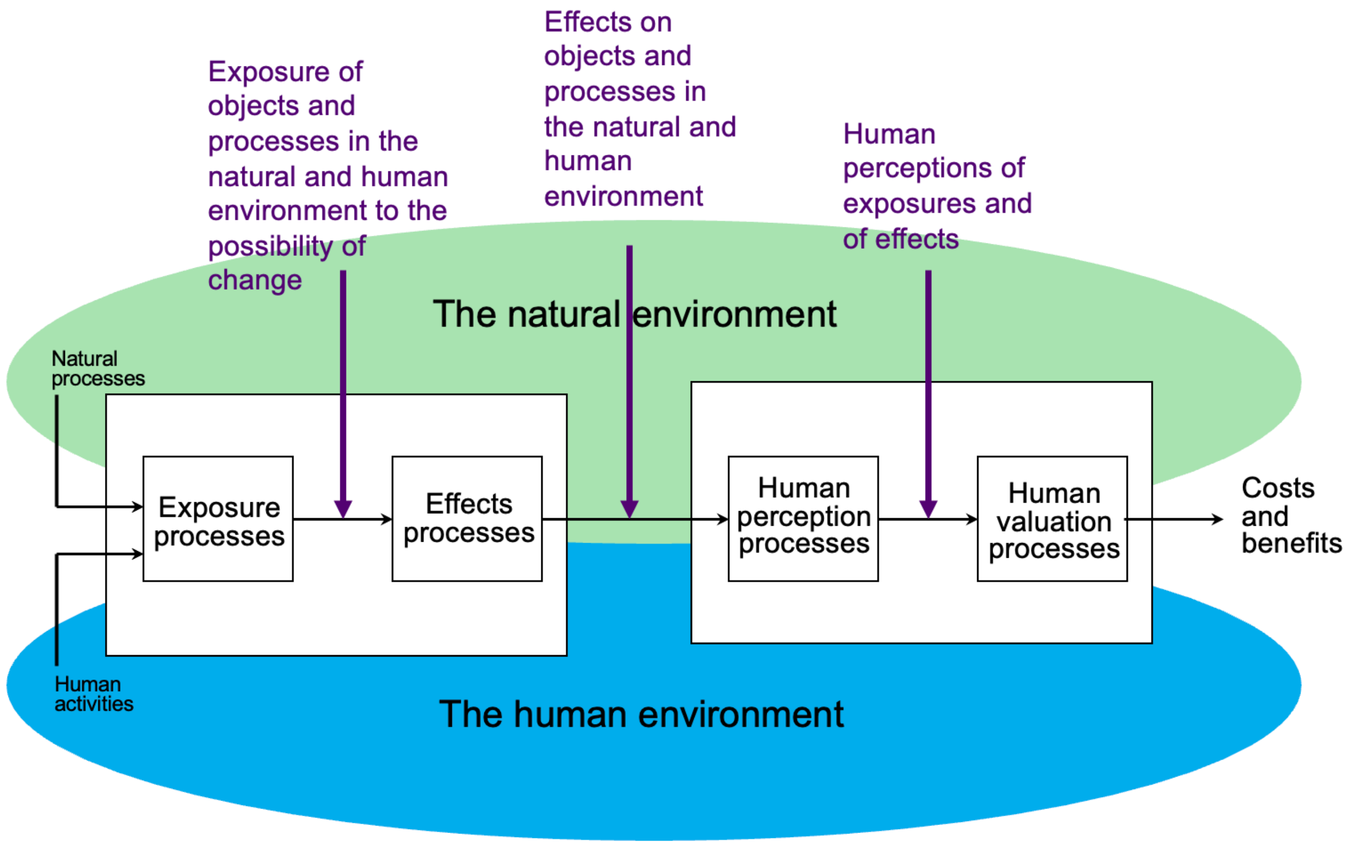

If you would like, later we can come back to talk more about these several parts

But, let's focus on just a few issues, starting with these first two boxes

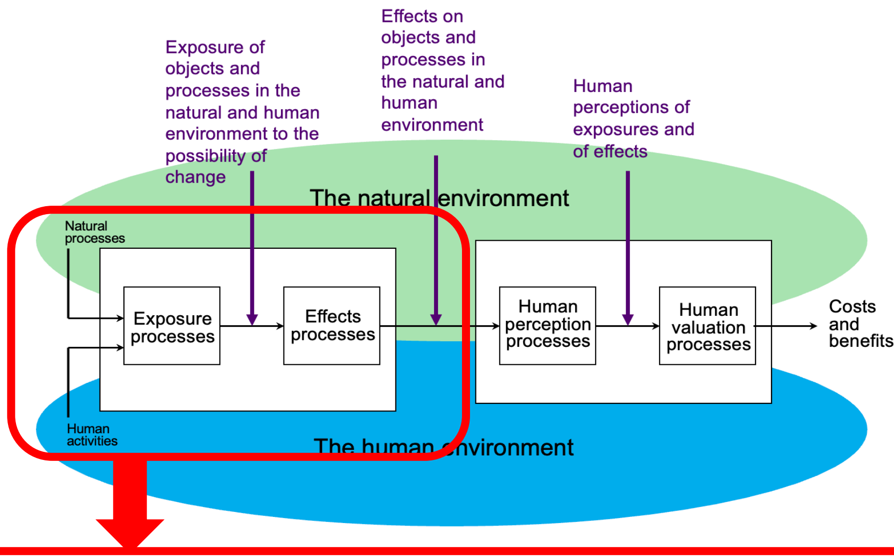

When there is uncertainty, out of the analysis we do in these first two boxes we may be able to obtain probability distributions for the effects of interest that might look something like this

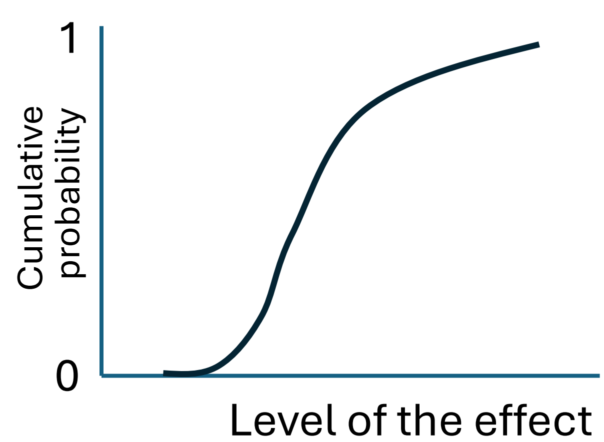

Where do we get the probabilities?

From available data whenever possible…

…from expert judgment when those data do not exist…

…and, if there aren't any real experts, and/or the specific evidence is so poor that eliciting experts is not a defensible strategy, using order of magnitude and bounding analysis.

The Neglected Art of Bounding Analysis

Environmental Science & Technology
April 1, 2001 / Volume 35 , Issue 7 / pp 162 A — 164 A.
M. GRANGER MORGAN

Are the answers provided by today's detailed risk analyses reasonable? Is valued insight being overlooked as a result of analysts' focus on the intimate details of environmental problems? If so, what can we do about this?

Environmental risk analysis has fallen into a standard front-to-back mode of operation: Estimate the magnitude and pattern of releases of the pollutants of concern; model their transport and transformation through the environment; estimate the location and physiological state of people, animals, and plants and the exposures they will receive; apply dose-response functions; and estimate the resulting impacts.

All of this makes perfect sense if the relevant science is pretty well known and good data are available on factors such as the behaviors of the populations at risk. However, in practice, the science is often highly uncertain. The release rates may not be known with precision. There is often great uncertainty about transport and transformation processes.

PERSPECTIVE

Use (and abuse) of expert elicitation in support of decision making for public policy

M. Granger Morgan[1]
Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213

Edited by William C. Clark, Harvard University, Cambridge, MA, and approved March 18, 2014 (received for review October 22, 2013)

The elicitation of scientific and technical judgments from experts, in the form of subjective probability distributions, can be a valuable addition to other forms of evidence in support of public policy decision making. This paper explores when it is sensible to perform such elicitation and how that can best be done. A number of key issues are discussed, including topics on which there are, and are not, experts who have knowledge that provides a basis for making informed predictive judgments; the inadequacy of only using qualitative uncertainty language; the role of cognitive heuristics and of overconfidence; the choice of experts; the development, refinement, and iterative testing of elicitation protocols that are designed to help experts to consider systematically all relevant knowledge when they make their judgments; the treatment of uncertainty about model functional form; diversity of expert opinion; and when it does or does not make sense to combine judgments from different experts. Although it may be tempting to view expert elicitation as a low-cost, low-effort alternative to conducting serious research and analysis, it is neither. Rather, expert elicitation should build on and use the best available research and analysis and be undertaken only when, given those, the state of knowledge will remain insufficient to support timely informed assessment and decision making.

Society often calls on experts for advice that requires judgments that go beyond well-established knowledge. In providing such judgments, it is common practice to use simulation models, engineering–economic assessment, and similar tools. Although such analytical strategies can provide valuable insight, they can never hope to include all relevant factors. In such situations, the community of applied decision analysis has long used quantitative expert judgments in the form of subjective probability distributions that have been elicited from relevant experts. Most such applications have been undertaken in support of decisions being made by private parties (1–4). Sometimes the resulting distributions are used directly, and sometimes they are fitted to formal functions and used in various Bayesian decision models (2, 5).

The use of expert elicitation in public sector decision making has been less common. Several studies have explored issues such as the health impacts of fine particle air pollution (6–12) and of lead pollution (13), the likely nature and extent of climate change (14–16), the various impacts that may result from climate change (17, 18), herbicide-tolerant oilseed crops (19), and the likely cost and performance of various energy technologies (20–24). The Environmental Protection Agency (EPA) has begun to make use of elicitation methods to address uncertain issues in environmental science (25), and those who work in both the Department of Energy and the Food and Drug Administration (FDA) have expressed interest in possibly using the method.

Done well, expert elicitation can make a valuable contribution to informed decision making. Done poorly it can lead to useless or even misleading results that lead decision makers astray, alienate experts, and wrongly discredit the entire approach. In what follows, I draw on relevant literature and 35 y of personal experience in designing and conducting substantively detailed expert elicitations, to suggest when it does and does not make sense to perform elicitations, how they should be designed and conducted, and how I believe the results should and should not be used. In contrast to much of the literature in Bayesian decision-making and applied decision analysis, my focus is on developing detailed descriptions of the state of understanding in some field of science or technology.

First, Are There Any Experts?

To conduct an expert elicitation, there must be experts whose knowledge can support informed judgment and prediction about the issues of interest. There are many topics about which people have extensive knowledge that provides little or no basis for making informed predictive judgments. For example, the further one moves away from questions whose answers involve matters of fact that are largely dependent on empirical natural or social science and well-validated models into realms in which individual and social behavior determine the outcomes of interest, the more one should ask whether expertise, with predictive capability, exists. For example, given a specified time series of future radiative forcing and other relevant physical variables, in my view, it is reasonable to ask climate scientists to make probabilistic judgments about average global temperature 150 y in the future. I am far less persuaded that it makes sense to ask "experts" questions that entail an assessment of how the stock market, or the price of natural gas will evolve over the next 25 y, or what the value of gross world product will be 150 y in the future.

The Interpretation of Probability

A subjectivist or Bayesian interpretation of probability (5, 26–28) is used when one makes subjective probabilistic assessments of the present or future value of uncertain quantities, the state of the world, or the nature of the processes that govern the world. In such situations, probability is viewed as a statement of an individual's belief, informed by all formal and informal evidence that he or she has available. Although subjective, such judgments cannot be arbitrary. They must conform to the laws of probability. Further, when large quantities of evidence are available on identical repeated events, one's subjective probability should converge to the classical frequentist interpretation of probability.

Partly as a result of their different training and professional cultures, different groups of experts display different views about the appropriateness of making subjective probabilistic judgments, and have different levels of willingness to make such judgments. Although every natural scientist and engineer I have ever interviewed seemed to think naturally in terms of subjective probabilities, others, such as some experts in the health sciences, have been far

Author contributions: M.G.M. designed research, performed research, and wrote the paper.
The author declares no conflict of interest.
This article is a PNAS Direct Submission.
[1]E-mail: granger.morgan@andrew.cmu.edu.
This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1319946111/-/DCSupplemental.

7176–7184 | PNAS | May 20, 2014 | vol. 111 | no. 20
www.pnas.org/cgi/doi/10.1073/pnas.1319946111

When uncertainty is high, and consequences could be serious, we should exercise caution

Journal of Risk Research **5** (4), 391–417 (2002)

Carfax Publishing
Taylor & Francis Group

Risk-based decision analysis in support of precautionary policies

MICHAEL L. DEKAY
Department of Engineering and Public Policy, Carnegie Mellon University (CMU) and H. John Heinz III School of Public Policy and Management, CMU, USA

MITCHELL J. SMALL*
Department of Engineering and Public Policy, and Department of Civil and Environmental Engineering, CMU, USA

PAUL S. FISCHBECK
Department of Engineering and Public Policy, and Department of Social and Decision Sciences, CMU, USA

R. SCOTT FARROW
Department of Engineering and Public Policy, and Center for the Study and Improvement in Regulation, CMU, USA

ALISON CULLEN
Daniel J. Evans School of Public Affairs and The National Research Center for Statistics and the Environment, University of Washington, USA

JOSEPH B. KADANE
Department of Statistics, CMU, USA

LESTER B. LAVE
Department of Engineering and Public Policy, and Graduate School of Industrial Administration, CMU, USA

M. GRANGER MORGAN
Department of Engineering and Public Policy, CMU, USA

KAZUHISA TAKEMURA
Institute of Policy and Planning Sciences, University of Tsukuba, Tsukuba, Japan

Abstract

A decision-analytic model for avoiding a risky activity is presented. The model considers the benefit and cost of avoiding the activity, the probability that the activity is unsafe, and scientific tests or studies that could be conducted to revise the probability that the activity is unsafe. For a single decision maker, thresholds are identified for his or her current subjective probability that the activity is unsafe. These thresholds indicate

whether the preferred course of action is avoiding the activity without further study, engaging in the activity without further study, or conducting a test or research programme to obtain additional information and following the result. When these thresholds are low, precautionary action is more likely to be warranted. When there are multiple stakeholders, differences in their perceptions of the benefit and cost of avoidance and differences in their perceptions of the accuracy of the additional information provided by the test or research programme combine to create differences in their decision thresholds. Thus, the model allows for the rational expression of differences among parties in a way that highlights disagreements and possible paths to conflict resolution. The model is illustrated with an application to phytosanitary standards in international trade and examined in terms of recent empirical research on lay perceptions of risks, benefits, and trust. Further research is suggested to improve the elicitation of model components, as a way of fostering the legitimate application of risk-based decision analysis in precautionary policy making.

KEY WORDS: decision analysis, precautionary principle, risk, standard of proof, threshold, trust

1. Introduction

Risk-based decision analysis provides a structured framework for making rational decisions when outcomes are uncertain. It does so by considering alternative actions intended to enhance benefits or avoid or mitigate losses, the possible outcomes associated with each of these actions, and the probabilities and relative desirabilities of these outcomes. It can also be used to evaluate options for reducing the uncertainty surrounding the probabilities and magnitudes of possible outcomes through diagnostic tests or additional research.

For many proposed activities, such as the introduction of a new drug, the administration of growth hormones to beef cattle, or the incineration of municipal or industrial waste, a clear set of benefits can be anticipated by those who support the action (though the value, allocation, and sustainability of these benefits may be subject to controversy). However, the risks associated with these programmes, should 'things not go as planned,' are usually much more uncertain. This is often the case when risks arise from low-probability, high-consequence events and processes within complex systems, such as those involving human health, ecology, and sociocultural institutions. The complexity and low probability of these effects can make it very difficult to assess and reduce the resulting uncertainty – so much so that some believe that problems of this type strain the capabilities of risk-based decision analysis.

Such concerns are part of the motivation for precautionary approaches that would substitute highly protective decision rules for the calculus and tradeoffs of risk-based decision making. In a classic paper, Page (1978) offered a characterization of the type of risk for which support for precautionary approaches is most likely to arise. Page described the (still) emerging set of technological, health, and environmental risks for which:

- there is an *ignorance of mechanism* – so that our knowledge of the physical processes that determine the likelihood and magnitude of the risk is poor

An example of overconfidence

Elicited cumulative distribution functions of the likely future costs of photovoltaic technology in 2030 obtained from seven energy experts in 2009-10.

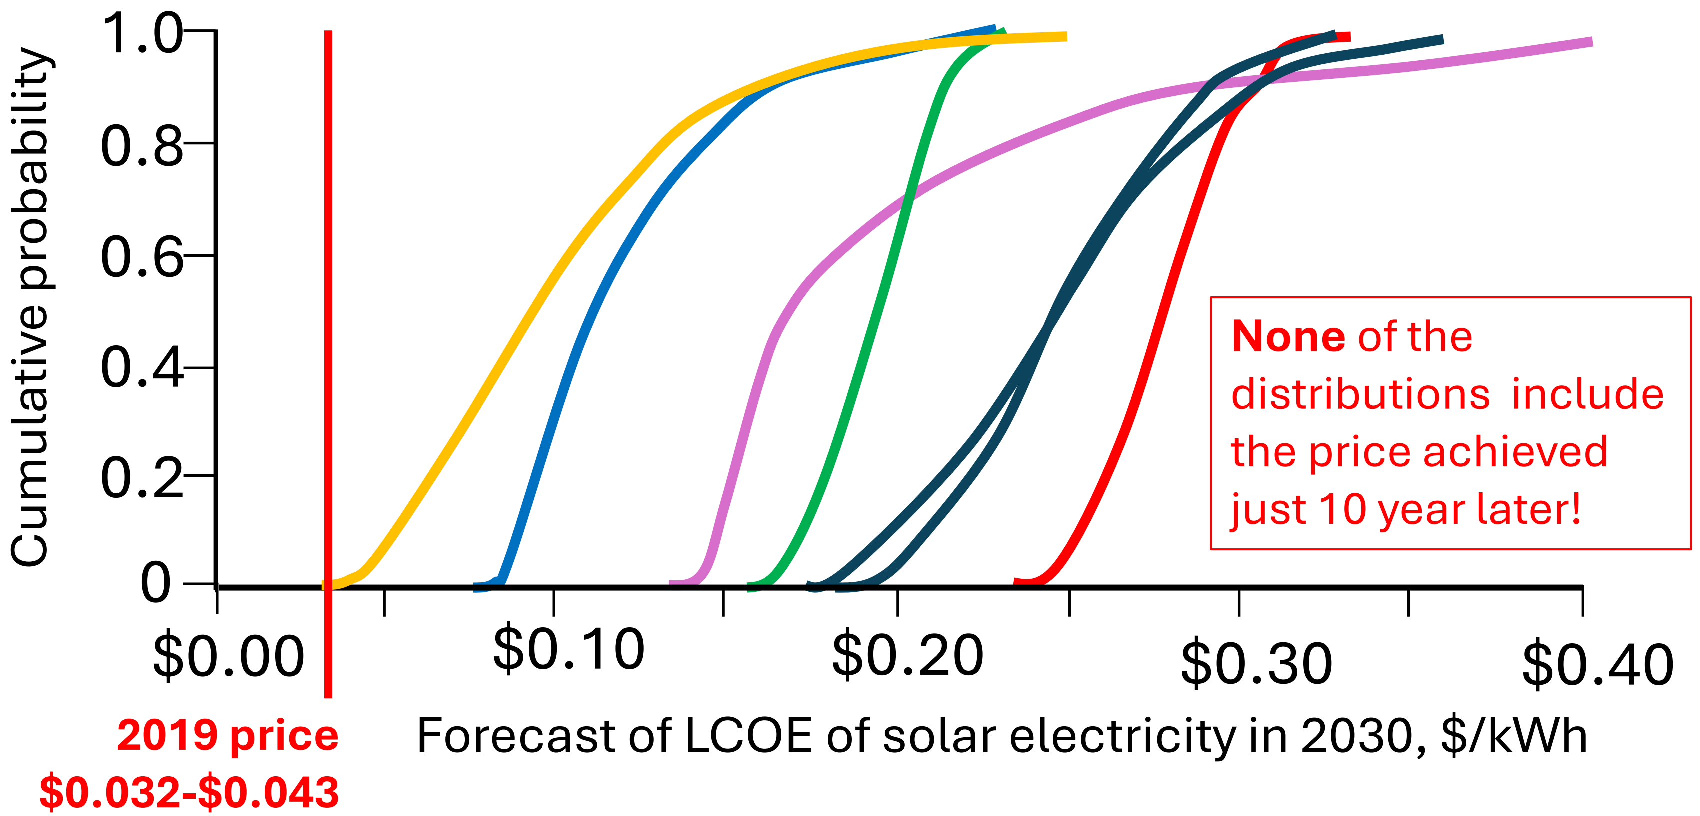

Image from: Savage, T., Davis, A., Fischhoff, B., & Morgan, M. G. (2021). A strategy to improve expert technology forecasts. *Proceedings of the National Academy of Sciences*, *118*(21), e2021558118.

Overconfidence is Ubiquitous

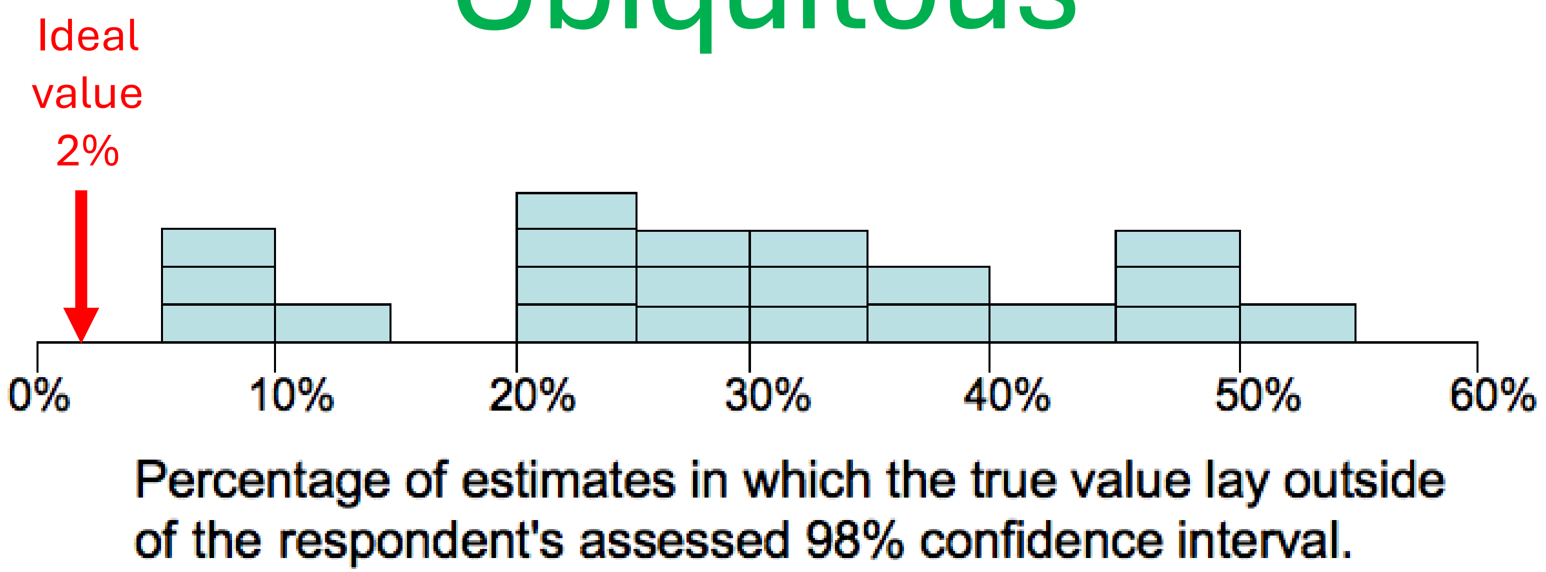

In 21 different studies involving over 10,000 assessment questions people were asked to assess 98% confidence intervals in a variety of judgments.

This histogram reports that value of the "surprise index," fractionn of the time answer lies outside the assessed 98% confidence interval (ideal value = 2%)

These results indicate clearly the ubiquitous tendency to overconfidence (i.e., assessed probabilities that are too narrow). A more detailed summary is provided in Morgan and Henrion, *Uncertainty*, Cambridge, 1990.

One more example for experts

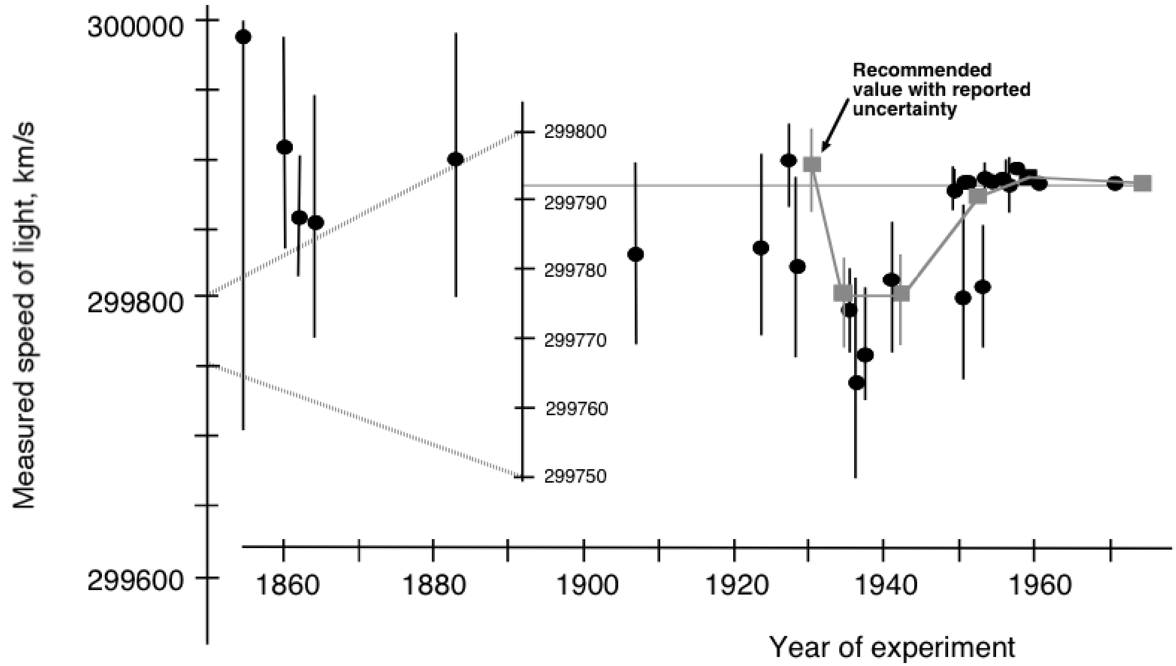

Published estimates of the speed of light. The light gray boxes that start in 1930 are the recommended values from the particle physics group that presumably include an effort to consider uncertainty arising from systematic error. Note that for over two decades the reported confidence intervals on these recommended values did not include the present best-measured value. Henrion and Fischhoff (1986), from which this figure is combined and redrawn, report that the same overconfidence is observed in the recommended values of a number of other physical constants.

A few issues involving the next two boxes

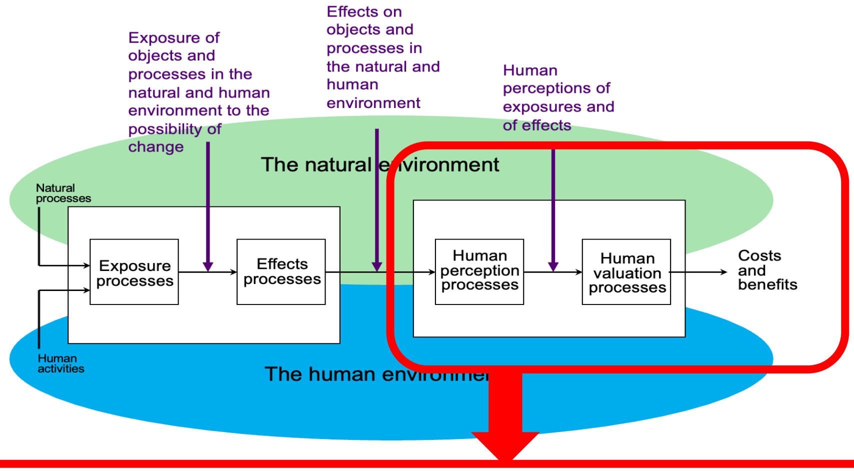

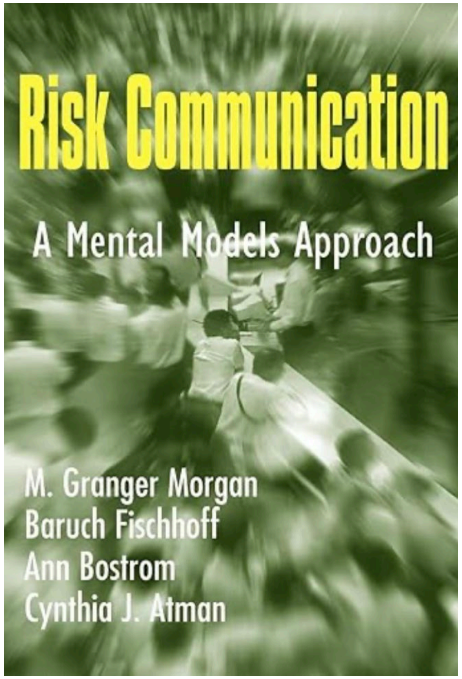

The sciences of science communication

Baruch Fischhoff[1]

Department of Engineering and Public Policy, and Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213-3890

Edited by Dietram A. Scheufele, University of Wisconsin, Madison, WI, and accepted by the Editorial Board May 23, 2013 (received for review September 19, 2012)

The May 2012 Sackler Colloquium on "The Science of Science Communication" brought together scientists with research to communicate and scientists whose research could facilitate that communication. The latter include decision scientists who can identify the scientific results that an audience needs to know, from among all of the scientific results that it would be nice to know; behavioral scientists who can design ways to convey those results and then evaluate the success of those attempts; and social scientists who can create the channels needed for trustworthy communications. This overview offers an introduction to these communication sciences and their roles in science-based communication programs.

decision making | environment | judgment | risk | uncertainty

We need science for many things, the most profound of which is supporting our sense of wonder about the world around us. Everyone has had moments in which science has solved or created a mystery. For me, one of those moments was as an undergraduate in Introductory Geology at Wayne State University, when Prof. (Joseph) Mozola explained the origins of the Harrisburg surface, which gives central Pennsylvania its distinctive lattice pattern of rivers and gaps. At that moment, I realized that all landforms, even the flats of Detroit, were created by processes that science could decode. Another such moment was reading Joanna Burger's *The Parrot Who Owns Me* (1) and getting some insight into how our white-crested cockatoo, Big Bird, appeared to have such human emotions, despite our species having diverged so long ago (even if he was born on Pittsburgh's South Side). A third was following a proposal from our son Ilya, an evolutionary biologist, to take out our lawn and let our yard go feral with native Pennsylvania flora—after which he taught us how to observe the seasons. Now, I can wonder why the first grackles arrived 10 days early this year, even though the main flock took over the yard on March 7th as usual. Was it a sign of climate change—or of my own improved powers of observation?

Science provides a sense of wonder not just from revealing the world, but also from showing that the world can be revealed. For that reason, I tell my advisees, in our Decision Science major, that they should join a research laboratory, any research laboratory, just to see how science is done. They may discover that they love the life, including the camaraderie of the graduate students and postdocs who are the engines of much science. Or, they may find the constant fussing over details to be "mind-numbingly dull" (to use the phrase that Calvin, of *Calvin and Hobbes*, applied to archaeology) and then choose to get their science from places like *Nova*, *Frontline*, the *New York Times*, and *Wired*. Either way, they will have seen the pursuit of uncertainty that distinguishes science from other forms of knowing and the scientific turn of mind: trying to get to the bottom of things, knowing that one never will.

Decisions That Require Understanding Science

Although people can choose not to do science, they cannot choose to ignore it. The products of science permeate their lives. Nuclear power, genetically modified organisms, nanotechnology, geoengineering, and xenotransplantation are just a few of today's realities that would have been impossible, even inconceivable, without scientific advances. Their very existence should be a source of wonder, even for people in whom they evoke a sense of terror ("we can do that?!").

The same mixed emotions may accompany phenomena that arise from the confluence of technology and society. How is it that human actions can both discover antibiotics and encourage microbial resistance to them, create "smart" electricity grids and threaten them with cyberattacks, and produce a food supply whose sheer abundance undermines the health of some people, while others still go hungry?

Without some grasp of the relevant science, it is hard to make informed decisions about these issues. Those include private decisions, such as whether to choose fuel-efficient vehicles, robot-guided surgery, or dairy products labeled as "produced by animals not treated with bovine growth hormone—not known to cause health effects." And they include public decisions, such as whether to support politicians who favor fuel efficiency, lobby for disclosing the risks of medical devices, or vote for referenda limiting agricultural biotechnology.

Effective science communications inform people about the benefits, risks, and other costs of their decisions, thereby allowing them to make sound choices. However, even the most effective communication cannot guarantee that people will agree about what those choices should be. Reasonable people may disagree about whether expecting X days less convalescence time is worth expecting Y% more complications, whether the benefits of organic produce justify Z% higher prices, or whether their community has been offered fair compensation for having a hazardous facility sited in its midst. Reasonable people may also disagree about what time frame to consider when making a decision (the next year? the next decade? the next century?), about whose welfare matters (their own? their family's? their country's? all humanity's?), or about whether to consider the decision-making process as well as its outcomes (are they being treated respectfully? are they ceding future rights?).

The goal of science communication is not agreement, but fewer, better disagreements. If that communication affords people a shared understanding of the facts, then they can focus on value issues, such as how much weight to give to nature, the poor, future generations, or procedural issues in a specific decision. To realize that potential, however, people need a venue for discussing value issues. Otherwise, those issues will percolate into discussions of science (2). For example, if health effects have legal standing in a facility-siting decision, but compensation does not, then local residents may challenge health effect studies, when their real concern is money. Science communication cannot succeed when people feel that attacking its message is the only way to get redress for other concerns (3).

Because science communication seeks to inform decision making, it must begin by listening to its audience, to identify the decisions that its members face—and, therefore, the information

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "The Science of Science Communication," held May 21–22, 2012, at the National Academy of Sciences in Washington, DC. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/science-communication.

Author contributions: B.F. wrote the paper.

The author declares no conflict of interest.

This article is a PNAS Direct Submission. D.A.S. is a guest editor invited by the Editorial Board.

[1]E-mail: baruch@cmu.edu.

www.pnas.org/cgi/doi/10.1073/pnas.1213273110

PNAS | August 20, 2013 | vol. 110 | suppl. 3 | 14033–14039

Downloaded from https://www.pnas.org by 71.112.157.17 on September 12, 2025 from IP address 71.112.157.17.

John D. Steinbruner

Before his untimely death, security expert John Steinbruner and colleagues…

…worked on managing the risks posed by synthetic biology. That work focused on biology that entailed risks that are ***much*** more immediate, and presumably larger, than those we are considering today.

BUT, one key contribution from their work was a set of arguments about the importance of being *open* and *transparent* about the research.

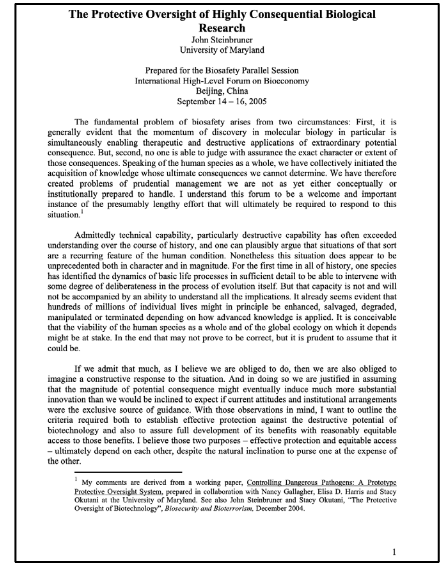

The Protective Oversight of Highly Consequential Biological Research

John Steinbruner
University of Maryland

Prepared for the Biosafety Parallel Session
International High-Level Forum on Bioeconomy
Beijing, China
September 14 – 16, 2005

The fundamental problem of biosafety arises from two circumstances: First, it is generally evident that the momentum of discovery in molecular biology in particular is simultaneously enabling therapeutic and destructive applications of extraordinary potential consequence. But, second, no one is able to judge with assurance the exact character or extent of those consequences. Speaking of the human species as a whole, we have collectively initiated the acquisition of knowledge whose ultimate consequences we cannot determine. We have therefore created problems of prudential management we are not as yet either conceptually or institutionally prepared to handle. I understand this forum to be a welcome and important instance of the presumably lengthy effort that will ultimately be required to respond to this situation.[1]

Admittedly technical capability, particularly destructive capability has often exceeded understanding over the course of history, and one can plausibly argue that situations of that sort are a recurring feature of the human condition. Nonetheless this situation does appear to be unprecedented both in character and in magnitude. For the first time in all of history, one species has identified the dynamics of basic life processes in sufficient detail to be able to intervene with some degree of deliberateness in the process of evolution itself. But that capacity is not and will not be accompanied by an ability to understand all the implications. It already seems evident that hundreds of millions of individual lives might in principle be enhanced, salvaged, degraded, manipulated or terminated depending on how advanced knowledge is applied. It is conceivable that the viability of the human species as a whole and of the global ecology on which it depends might be at stake. In the end that may not prove to be correct, but it is prudent to assume that it could be.

If we admit that much, as I believe we are obliged to do, then we are also obliged to imagine a constructive response to the situation. And in doing so we are justified in assuming that the magnitude of potential consequence might eventually induce much more substantial innovation than we would be inclined to expect if current attitudes and institutional arrangements were the exclusive source of guidance. With those observations in mind, I want to outline the criteria required both to establish effective protection against the destructive potential of biotechnology and also to assure full development of its benefits with reasonably equitable access to those benefits. I believe those two purposes – effective protection and equitable access – ultimately depend on each other, despite the natural inclination to purse one at the expense of the other.

[1] My comments are derived from a working paper, Controlling Dangerous Pathogens: A Prototype Protective Oversight System, prepared in collaboration with Nancy Gallagher, Elisa D. Harris and Stacy Okutani at the University of Maryland. See also John Steinbruner and Stacy Okutani, "The Protective Oversight of Biotechnology", *Biosecurity and Bioterrorism*, December 2004.

1

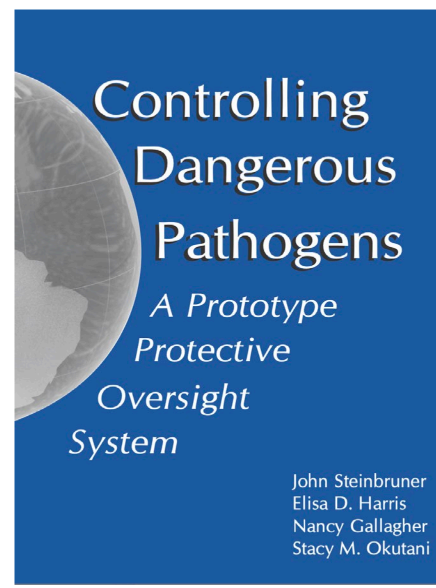

Source: UNIVERSITY OF MARYLAND | DRUM
DIGITAL REPOSITORY AT THE UNIVERSITY OF MARYLAND

This suggests the need to consider this balance

Image source: Data Force

With that I’ll say thanks.

I’ll be happy to talk more about any of this during the discussion.