

Communicating Uncertainty Fulfilling the Duty to Inform

Scientists are often hesitant to share their uncertainty with decisionmakers who need to know it. With an understanding of the reasons for their reluctance, decisionmakers can create the conditions needed to facilitate better communication.

Experts' knowledge has little practical value unless its recipients know how sound it is. It may even have negative value, if it induces unwarranted confidence or is so hesitant that other, overstated claims push it aside. As a result, experts owe decisionmakers candid assessments of what they do and do not know.

Decisionmakers ignore those assessments of confidence (and uncertainty) at their own peril. They may choose to act with exaggerated confidence, in order to carry the day in political debates, or with exaggerated hesitancy, in order to avoid responsibility. However, when making decisions, they need to know how firm the ground is under their science, lest opponents attack unsuspected weaknesses or seize initiatives that might have been theirs.


Sherman Kent, the father of intelligence analysis, captured these risks in his classic essay "Words of Estimative Probability," showing how leaders can be misled by ambiguous ex-

pressions of uncertainty. As a case in point, he takes a forecast from National Intelligence Estimate 29-51, *Probability of an Invasion of Yugoslavia in 1951*: "Although it is impossible to determine which course the Kremlin is likely to adopt, we believe that the extent of Satellite military and propaganda preparations indicates that an attack on Yugoslavia in 1951 should be considered a serious possibility." When he asked other members of the Board of National Estimates "what odds they had had in mind when they agreed to that wording," their answers ranged from 1:4 to 4:1. Political or military leaders who interpret that forecast differently might take very different actions, as might leaders who make different assumptions about how much the analysts agree.

Figure 1 shows the same problem in a very different domain. In fall 2005, as avian flu loomed, epidemiologist Larry Brilliant convened a meeting of public health experts, able to assess the threat, and technology experts, able to assess the options for keeping society going, if worst came to worst. In preparation for the meeting, we surveyed their beliefs. The

FIGURE 1

Judgments of “the probability that H5N1 will become an efficient human-to-human transmitter (capable of being propagated through at least two epidemiological generations of humans) some time during the next 3 years.”



Data collected in October 2005. [Source: W. Bruine de Bruin, B. Fischhoff, L. Brilliant, and D. Caruso, “Expert Judgments of Pandemic Influenza,” *Global Public Health* 1, no. 2 (2006): 178–193]

figure shows their answers to the first question on our survey, eliciting the probability that the virus would become an efficient human-to-human transmitter in the next three years.

The public health experts generally saw a probability around 10%, with a minority seeing a higher one. The technology experts, as smart a lay audience as one could imagine, saw higher probabilities. Possibly, they had heard both groups of public health experts and had sided with the more worried one. More likely, though, they had seen the experts’ great concern and then assumed a high probability. However, with an anticipated case-fatality rate of 15% (provided in response to another question), a 10% chance of efficient transmission is plenty of reason for concern.

Knowing that probability is essential to orderly decision-making. However, it is nowhere to be found in the voluminous coverage of H5N1. Knowing that probability is also essential to evaluating public health officials’ performance. If they seemed to be implying a 70% chance of a pandemic, then they may seem to have been alarmist, given that none occurred. That feeling would have been reinforced if they seemed equally alarmed during the H1N1 mobilization, unless they said clearly that they perceived a low-probability event with very high consequences should it come to pass.

For an expert who saw a 10% chance, a pandemic would have been surprising. For an expert who saw a 70% chance, the absence of a pandemic would have been. However, neither surprise would render the prediction indefensible. That happens only when an impossible event (probability = 0%) occurs or a sure thing (probability = 100%) does not. Thus,

for an individual event, any probability other than 0% or 100% is something of a hedge, because it cannot be invalidated. For a set of events, though, probability judgments should be calibrated, in the sense that events with a 70% chance happen 70% of the time, 10% events 10% of the time, and so on.

After many years of eliciting probability judgments from foreign policy experts, then seeing what happened, Philip Tetlock concluded that they were generally overconfident. That is, events often did not happen when the experts were confident that they would or happened when the experts were confident that they would not. In contrast, Allan Murphy and Robert Winkler found that probability-of-precipitation (PoP) forecasts express appropriate confidence. It rains 70% of the time that they give it a 70% chance.

Thus, decisionmakers need to know not only how confident their experts are, but also how to interpret those expressions of confidence. Are they like Tetlock’s experts, whose confidence should be discounted, or like Murphy and Winkler’s, who know how much they know? Without explicit statements of confidence, evaluated in the light of experience, there is no way of knowing.

Kent ended his essay by lamenting analysts’ reluctance to be that explicit, a situation that largely persists today. And not just in intelligence. For example, climate scientists have wrangled with similar resistance, despite being burned by forecasts whose imprecision allowed them to be interpreted as expressing more confidence than was actually intended. Even PoP forecasts face sporadic recidivism from weather

forecasters who would rather communicate less about how much they know.

Reasons for reluctance

Experts' reluctance to express their uncertainty has understandable causes, and identifying those causes can help us develop techniques to make expert advice more useful.

Experts see uncertainty as misplaced imprecision. Some experts feel that being explicit about their uncertainties is wrong, because it conveys more precision than is warranted. In such cases, the role of uncertainty in decisionmaking has not been explained well. Unless decisionmakers are told how strong or shaky experts' evidence is, then they must guess. If they guess wrong, then the experts have failed them, leading to decisions made with too much or too little confidence and without understanding the sources of their uncertainty. Knowing about those sources allows decisionmakers to protect themselves against vulnerabilities, commission the analyses needed to reduce them, and design effective actions. As a result, decisionmakers have a right, and a responsibility, to demand that uncertainty be disclosed.

In other cases, experts are willing to share their uncertainty but see no need because it seems to go without saying: Why waste decisionmakers' valuable time by stating the obvious? Such reticence reflects the normal human tendency to exaggerate how much of one's knowledge goes without saying. Much of cognitive social psychology's stock in trade comes from documenting variants of that tendency. For example, the "common knowledge effect" is the tendency to believe that others share one's beliefs. It creates unpleasant surprises when others fail to read between the lines of one's unwittingly incomplete messages. The "false consensus effect" is the tendency to believe that others share one's attitudes. It creates unpleasant surprises when those others make different choices, even when they see the same facts, because they have different objectives.

Because we cannot read one another's minds, all serious communication needs some empirical evaluation, lest its effectiveness be exaggerated. That is especially true when communicating about unusual topics to unfamiliar audiences, where communicators can neither assume shared understanding nor guess how they might be misinterpreted. That is typically the lot of experts charged with conveying their knowledge and its uncertainties to decisionmakers. As a result, they bear a special responsibility to test their messages before transmission. At the least, they can run them by some nonexperts somewhat like the target audience, asking them to paraphrase its content, "just be sure that it was clear."

When decisionmakers receive personal briefings, they

can ask clarifying questions or make inferences that reveal how well they have mastered the experts' message. If decisionmakers receive a communication directed at a general audience, though, they can only hope that its senders have shown due diligence in disambiguating the message. If not, then they need to add a layer of uncertainty, as a result of having to guess what they might be missing in what the experts are trying to say. In such cases, they are not getting full value from the experts' work.

Experts do not expect uncertainties to be understood. Some experts hesitate to communicate uncertainties because they do not expect nonexpert audiences to benefit. Sometimes, that skepticism comes from overestimating how much they need to communicate in order to bring value. Decisionmakers need not become experts in order to benefit from knowing about uncertainties. Rather, they can often derive great marginal utility from authoritative accounts conveying the gist of the key issues, allowing them to select those where they need to know more.

In other cases, though, experts' reluctance to communicate uncertainties comes from fearing that their audience would not understand them. They might think that decisionmakers lack the needed cognitive abilities and substantive background or that they cannot handle the truth. Therefore, the experts assume that decisionmakers need the sureties of uncertainty-free forecasts. Such skepticism takes somewhat different forms with the two expressions of uncertainty: summary judgments of its extent (for example, credible intervals around possible values) and analyses of its sources (for example, threats to the validity of theories and observations).

With summary judgments, skeptical experts fear that nonexperts lack the numeracy needed to make sense of probabilities. There are, indeed, tests of numeracy that many laypeople fail. However, these tests typically involve abstract calculations without the context that can give people a feeling for the numbers. Experts can hardly support the claim that lay decisionmakers are so incompetent that they should be denied information relevant to their own well-being.

The expressions of uncertainty in Figure 1 not only have explicit numeric probabilities but are also attached to an event specified precisely enough that one can, eventually, tell whether it has happened. That standard is violated in many communications, which attach vague verbal quantifiers (such as rare, likely, or most) to vague events (such as environmental damage, better health, or economic insecurity). Unless the experts realize the ambiguity in their communications, they may blame their audience for not "getting" messages that it had little chance to understand.

TABLE 1**Brier Score Decomposition for evaluating probabilistic forecasts**

Three kinds of performance:

- **Knowledge.** Forecasts are better the more often they come true (for example, it rains when rain is forecasted and not when it is not). However, that kind of accuracy does not tell the whole story, without considering how hard the task is (for example, forecasts in the Willamette Valley would often be correct just predicting rain in the winter and none in the summer).
- **Resolution.** Assessing uncertainty requires discriminating among different states of knowledge. It is calculated as the variance in the percentage of correct predictions associated with different levels of expressed confidence.
- **Calibration.** Making such discriminations useful to decisionmakers requires conveying the knowledge associated with each level. Perfect calibration means being correct XX% of the time when one is XX% confident. Calibration scores reflect the squared difference between those two percentages, penalizing those who are especially over- or underconfident. [See A. H. Murphy, "A New Vector Partition of the Probability Score," *Journal of Applied Meteorology* 12 (1973): 595–600.]

TABLE 2**Boundary conditions on experimental tasks studying decisionmaking performance**

- The tasks are clearly described. That can produce better decisions, if it removes the clutter of everyday life, or worse decisions, if that clutter provides better context, such as what choices other people are making.
- The tasks have low stakes. That can produce better decisions, if it reduces stress, or worse decisions, if it reduces motivation.
- The tasks are approved by university ethics committees. That can produce better decisions, if it reduces worry about being deceived, or worse decisions, if it induces artificiality.
- The tasks focus on researchers' hypotheses. That can produce better decisions, if researchers are looking for decisionmakers' insights, or worse decisions, if they are studying biases.

[Adapted from B. Fischhoff and J. Kadvan, *Risk: A Very Short Introduction* (Oxford: Oxford University Press, 2011), p. 110]

PoP forecasts are sporadically alleged to confuse the public. The problem, though, seems to lie with the event, not the probability attached to it. People have a feeling for what 70% means but are often unsure whether it refers to the fraction of the forecast period that it will rain, the fraction of the area that it will cover, or the chance of a measurable amount at the weather station. (In the United States, it is the last.) Elite decisionmakers may be able to demand clarity and relevance from the experts who serve them. Members of the general public have little way to defend themselves against ambiguous communications or against accusations of having failed to comprehend the incomprehensible.

Analogous fears about lay incompetence underlie some experts' reluctance to explain their uncertainty. As with many intuitions about others' behavior, these fears have some basis in reality. What expert has not observed nonexperts make egregious misstatements about essential scientific facts? What expert has not heard, or expressed, concerns about the decline of STEM (science, technology, engineering, and mathematics) education and literacy? Here too, intuitions can be misplaced, unless supported by formal analysis of what people need to know and empirical study of what they already do know.

Decisionmakers need to know the facts (and attendant uncertainties) relevant to the decisions they face. It is nice to know many other things, especially if such background knowledge facilitates absorbing decision-relevant information. However, one does not need coursework in ecology or biochemistry to grasp the gist of the uncertainty captured in statements such as, "We have little experience with the effects of current ocean acidification (or saline runoff from hydrofracking or the effects of large-scale wind farms on developing-country electrical grids, or the off-label use of newly approved pharmaceuticals)."

Experts anticipate being criticized for communicating uncertainty. Some experts hesitate to express uncertainties because they see disincentives for such candor. Good decisionmakers want to know the truth, however bad and uncertain. That knowledge allows them to know what risks they are taking, to prepare for surprises, and to present their choices with confidence or caution. However, unless experts are confident that they are reporting to good decisionmakers, they need assurance that they will be protected if they report uncertainties.

Aligning the incentives of experts and decisionmakers is a basic challenge for any organization, analyzed in the recent National Research Council report *Intelligence Analysis for Tomorrow*, sponsored by the Office of the Director of National Intelligence. In addition to evaluating approaches to



MARINA ZURKOW, *Slurb*, Video animation; 17:42 minutes, 2009.
Courtesy of Bryce Wolkowitz Gallery, New York.

analysis, the report considers the organizational behavior research relevant to recruiting, rewarding, and retaining analysts capable of using those approaches, and their natural talents, to the fullest. Among other things, it recommends having analysts routinely assign probabilities for their forecasts precisely enough to be evaluated in the light of subsequent experience. Table 1 describes distinctions among forms of expert performance revealed in probabilistic forecasts.

Without such policies, experts may realistically fear that decisionmakers will reward bravado or waffling over candor. Those fears can permeate work life far beyond the moments of truth where analyses are completed and communicated. Experts are people, too, and subject to rivalries and miscommunication among individuals and groups, even when working for the same cause. Requiring accurate assessment of uncertainty reduces the temptation to respond strategically rather than honestly. Conversely, it protects sincere experts from the demoralization that comes with seeing others work the system and prosper.

Even when their organization provides proper incentives, some experts might fear their colleagues' censure should their full disclosure of uncertainty reveal their field's "trade secrets." Table 2 shows some boundary conditions on results from the experiments that underpin much decision-making research, cast in terms of how features of those studies affect the quality of the performance that they reveal. Knowing them is essential to applying that science in the right places and with the appropriate confidence. However, declaring them acknowledges limits to science that has of-

ten had to fight for recognition against more formal analyses (such as economics or operations research), even though the latter necessarily neglect phenomena that are not readily quantified. Decisionmakers need equal disclosure from all disciplines, lest they be unduly influenced by those that oversell their wares.

Experts do not know how to express their uncertainties.

A final barrier faces some experts who realize the value of assessing uncertainty, trust decisionmakers to understand well-formulated communications, and expect to be rewarded for doing so: They are uncertain how to perform those tasks to a professional standard. Although all disciplines train practitioners to examine their evidence and analyses critically, not all provide training in summarizing their residual uncertainties in succinct, standard form. Indeed, some disciplines offer only rudimentary statistical training for summarizing the variability in observations, which is one input to overall uncertainty.

For example, the complexity of medical research requires such specialization that subject-matter experts might learn just enough to communicate with the experts in a project's "statistical core," entrusted with knowing the full suite of statistical theory and methods. Although perhaps reasonable on other grounds, that division of labor can mean that subject-matter experts understand little more than rudimentary statistical measures such as P values. Moreover, they may have limited appreciation of concepts, such as how statistical significance differs from practical significance, assumes representative sampling, and depends on both

FIGURE 2

Explanation of uncertainty terms in intelligence analysis

What we mean when we say: An explanation of estimative language

When we use words such as “we judge” or “we assess”—terms we use synonymously—as well as “we estimate,” “likely” or “indicate,” we are trying to convey an analytical assessment or judgment. These assessments, which are based on incomplete or at times fragmentary information are not a fact, proof, or knowledge. Some analytical judgments are based directly on collected information; others rest on previous judgments, which serve as building blocks. In either type of judgment, we do not have “evidence” that shows something to be a fact or that definitively links two items or issues.

Intelligence judgments pertaining to likelihood are intended to reflect the Community’s sense of the probability of a development or event. Assigning precise numerical ratings to such judgments would imply more rigor than we intend. The chart below provides a rough idea of the relationship of terms to each other.

Remote	Unlikely	Even chance	Probably, likely	Almost certainly
--------	----------	-------------	------------------	------------------

We do not intend the term “unlikely” to imply an event will not happen. We use “probably” and “likely” to indicate there is a greater than even chance. We use words such as “we cannot dismiss,” “we cannot rule out,” and “we cannot discount” to reflect an unlikely—or even remote—event whose consequences are such it warrants mentioning. Words such as “may be” and “suggest” are used to reflect situations in which we are unable to assess the likelihood generally because relevant information is nonexistence, sketchy, or fragmented.

[Source: Office of the Director for National Intelligence, *Prospects for Iraq’s Stability: A Challenging Road Ahead* (Washington, DC: 2007)]

sample size and measurement precision.

Figure 2, taken from the 2007 National Intelligence Estimate for Iraq, shows an important effort by U.S. analysts to clarify the uncertainties in their work. It stays close to its experts’ intuitive ways of expressing themselves, requiring it to try to explain those terms to decisionmakers. Figure 3 represents an alternative strategy, requiring experts to take a step toward decisionmakers and summarize their uncertainties in explicit decision-relevant terms. Namely, it asks

experts for minimum and maximum possible values, along with some intermediate ones. As with PoP forecasts, sets of such judgments can be calibrated in the light of experience. For example, actual values should turn out to be higher than the 0.95 fractile only 5% of the time. Experts are overconfident if more than 10% of actual events fall outside the 0.05 to 0.95 range; they are underconfident if there are too few such surprises.

Some experts object in principle to making such judgments, arguing that all probabilities should be “objective” estimates of the relative frequency of identical repeated events (such as precipitation). Advocates of the alternative “subjectivist” principle argue that even calculated probabilities require judgments (such as deciding that climate conditions are stable enough to calculate meaningful rates). Whatever the merits of these competing principles, as a practical matter, decisionmakers need numeric probabilities (or verbal ones with clear, consensual numeric equivalents). Indeed, subjective probabilities are formally defined (in utility theory) in practical terms; namely, the gambles that people will take based on them. Experts are only human if they prefer to express themselves in verbal terms, even when they receive quantitative information. However, doing so places their preferences above decisionmakers’ needs.

The decisionmakers’ role

Just as decisionmakers need help from experts in order to make informed choices, experts need help from decisionmakers in order to provide the most useful information. Here are steps that decisionmakers can take for overcoming each source of experts’ reluctance to express uncertainty. Implementing them is easier for elite decisionmakers, who can make direct demands on experts, than for members of the general public, who can only hope for better service.

If experts see little value in expressing their uncertainty, show how decisions depend on it. Assume that experts cannot guess the nature of one’s decisions or one’s current thinking about them. Experts need direction, regarding all three elements of any decision: the decisionmakers’ goals, options, and beliefs (about the chances of achieving each goal with each option). The applied science of translating decisionmakers’ perceptions into formal terms is called decision analysis. Done well, it helps decisionmakers clarify their thinking and their need to understand the uncertainty in expert knowledge. However, even ordinary conversations can reveal much about decisionmakers’ information needs: Are they worried about that? Don’t they see those risks? Aren’t they considering that option? Decisionmakers must require those interactions with experts.

Whether done directly or through intermediaries such as survey researchers or decision analysts, those interactions should follow the same pattern. Begin by having decisionmakers describe their decisions in their own natural terms, so that any issue can emerge. Have them elaborate on whatever issues they raise, in order to hear them out. Then probe for other issues that experts expected to hear, seeing if those issues were missed or the experts were mistaken. Making decisions explicit should clarify the role of uncertainty in them.

If experts fear being misunderstood, insist that they trust their audience. Without empirical evidence, experts cannot be expected to know what decisionmakers currently know about a topic or could know if provided with well-designed communications. Human behavior is too complex for even behavioral scientists to make confident predictions. As a result, they should discipline their predictions with data, avoiding the sweeping generalizations found in popular accounts (“people are driven by their emotions,” “people are overconfident,” “people can trust their intuitions”).

Moreover, whatever their skepticism about the public’s abilities, experts often have a duty to inform. They cannot, in Bertolt Brecht’s terms, decide that the public has “forfeited [their] confidence and could only win it back by redoubled labor.” Even if experts might like a process that “dissolved the people and elected another one,” earning the public’s trust requires demonstrating their commitment to its well-being and their competence in their work. One part of that demonstration is assessing and communicating their uncertainty.

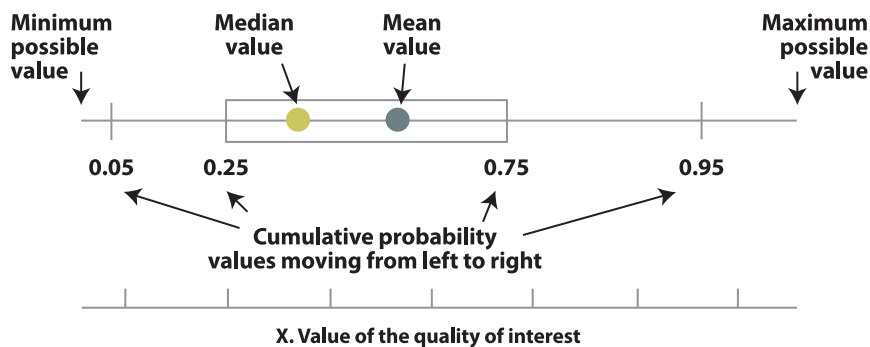
If experts anticipate being punished for candor, stand by them. Decisionmakers who need to know about uncertainty must protect those who provide it. Formal protection requires personnel policies that reward properly calibrated expressions of uncertainty, not overstated or evasive analyses. It also requires policies that create the feedback needed for learning and critical internal discussion. It requires using experts to inform decisions, not to justify them.

Informal protection requires decisionmakers to demonstrate a commitment to avoid hindsight bias, which understates uncertainty in order to blame experts for not having made difficult predictions, and hindsight bias, which overstates uncertainty in order to avoid blame by claiming that no one could have predicted the unhappy events following decisionmakers’ choices. Clearly explicating uncertainty conveys that commitment. It might even embolden experts to defy the pressures arising from professional norms and interdisciplinary competition to understate uncertainty.

If experts are unsure how to express themselves, provide standard means. At least since the early 18th century, when Daniel Bernoulli introduced his ideas on probability, scientists have struggled to conceptualize chance and uncertainty. There have been many thoughtful proposals for formalizing those concepts, but few that have passed both the theoretical test of rigorous peer review and the practical test of applicability. Decisionmakers should insist that experts use those proven methods for expressing uncertainty. Whatever its intuitive appeal, a new method is likely to share flaws with some of the other thoughtful approaches that have

FIGURE 3

Recommended box plot for expressing uncertainty



Source: P. Campbell, “Understanding the Receivers and the Receptions of Science’s Uncertain Messages,” *Philosophical Transactions of the Royal Society* 369 (2011): 4891–4912.

TABLE 3**FAQ for experts worried about providing subjective probability judgments**

Concern 1: People will misinterpret them, inferring greater precision than I intended.

Response: Behavioral research has found that most people like receiving explicit quantitative expressions of uncertainty (such as credible intervals), can interpret them well enough to extract their main message, and misinterpret verbal expressions of uncertainty (such as “good” evidence or “rare” side effect). For most audiences, misunderstanding is more likely with verbal expressions.

Concern 2: People cannot use probabilities.

Response: Behavioral research has found that even laypeople can provide high-quality probability judgments, if they are asked clear questions and given the chance to reflect on them. That research measures the quality of those judgments in terms of their internal consistency (or coherence) and external accuracy (or calibration).

Concern 3: My credible intervals will be used unfairly in performance evaluations.

Response: Probability judgments can protect experts by having them express the extent of their knowledge, so that they are not unfairly accused of being too confident or not confident enough.

fallen by the wayside. Table 3 answers some questions experts might have.

Decisionmakers should then meet the experts halfway by mastering the basic concepts underlying those standard methods. The National Research Council report mentioned earlier identified several such approaches whose perspective should be familiar to any producer or consumer of intelligence analyses. For example, when decisionmakers know that their opponents are trying to anticipate their choices, they should realize that game theory addresses such situations and therefore commission analyses where needed and know how far to trust their conclusions. A companion volume, *Intelligence Analysis: Behavioral and Social Science Foundations*, provides elementary introductions to standard methods, written for decisionmakers.

Experts are sometimes reluctant to provide succinct accounts of the uncertainties surrounding their work. They

may not realize the value of that information to decisionmakers. They may not trust decisionmakers to grasp those accounts. They may not expect such candor to be rewarded. They may not know how to express themselves.

Decisionmakers with a need, and perhaps a right, to know about uncertainty have ways, and perhaps a responsibility, to overcome that reluctance. They should make their decisions clear, require such accounts, reward experts who provide them, and adopt standard modes of expression. If successful, they will make the critical thinking that is part of experts’ normal work accessible to those who can use it.

Recommended reading

- B. Fischhoff, N. Brewer, and J. S. Downs, eds., *Communicating Risks and Benefits: An Evidence-Based User’s Guide* (Washington, DC: U.S. Food and Drug Administration, 2011).
- B. Fischhoff and C. Chauvin, eds., *Intelligence Analysis: Behavioral and Social Science Foundations* (Washington, DC: National Academies Press, 2011).
- B. Fischhoff and J. Kadvany, *Risk: A Very Short Introduction* (Oxford: Oxford University Press, 2011).
- R. Gregory, L. Failing, M. Harstone, G. Long, T. McDaniels, and D. Ohlson, *Structured Decisionmaking* (Chichester, UK: Wiley-Blackwell, 2012).
- M. G. Morgan, B. Fischhoff, A. Bostrom, and C. Atman, *Risk Communication: A Mental Models Approach* (New York: Cambridge University Press, 2002).
- A. H. Murphy and R. L. Winkler, “Probability Forecasting in Meteorology,” *Journal of the American Statistical Association* 79 (1984): 489–500.
- National Research Council, *Intelligence Analysis for Tomorrow* (Washington, DC: National Academies Press, 2011).
- R. A. Nickerson, “How We Know—and Sometimes Miscalculate—What Others Know: Imputing Our Own Knowledge to Others,” *Psychological Bulletin* 125 (1999): 737–759.
- A. O’Hagan et al., *Uncertain Judgments: Eliciting Expert Probabilities* (Chichester, UK: Wiley, 2006).
- P. E. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton, NJ: Princeton University Press, 2006).

Baruch Fischhoff (baruch@cmu.edu) is the Howard Heinz University Professor in the Departments of Social and Decision Sciences and of Engineering and Public Policy at Carnegie Mellon University, where he heads the Decision Science major.