Managing Knowledge to Reduce Groupthink

Baruch Fischhoff, PhD
Howard Heinz University Professor
Department of Social and Decision Sciences
Department of Engineering and Public Policy
Carnegie Mellon University
baruch@cmu.edu

Prepared for Unclassified Briefing Session on

Groupthink and Intelligence Analysis:
Thinking in Groups without Groupthink

Organized by
Office of Analytic Integrity and Standards
Office of the Director of National Intelligence

Monday, September 25, 2006
12:00-3:30pm
The Capitol, Room HC-7

Managing Knowledge to Reduce Groupthink

I am honored to be here. Thank you for the opportunity to speak about this critical problem.

I am Baruch Fischhoff, Howard Heinz University Professor at Carnegie Mellon University. I am a cognitive psychologist by training. My research focuses on helping people to deal with complex risks. I am a member of the Institute of Medicine of the National Academy of Sciences, the Department of Homeland Security's Science and Technology Advisory Committee, and the Environmental Protection Agency's Science Advisory Board, where I chair the Homeland Security Advisory Committee. I am a past President of the Society for Risk Analysis and of the Society for Judgment and Decision Making.

I have had the opportunity to interact with members of the intelligence community, at different times during my professional career, including participation in the DARPA-sponsored program on decision analysis, which I joined after getting my PhD at the Hebrew University of Jerusalem, in 1974. In that period, the CIA's Dick Heuer produced his book, The Psychology of Intelligence Analysis. Recently revised, it still bears reading, as a translation of behavioral research into practical terms.

I lived in Israel during the uneasy peace that preceded the 1973 (Yom Kippur) War and the recriminations that followed it. A common view was that Israel had been the victim of groupthink, and had paid a heavy price for it. The view was a basis for the Agranath Commission's recommendation of setting up a parallel, independent analytical authority in the Foreign Ministry.

At the time, I was completing my dissertation – on hindsight bias. One of the glories of being totally immersed in a dissertation is seeing its implications everywhere. I saw them in the second guessing of Israel's leaders that was going on around me – and in Irving Janis's book, Victims of Groupthink, which was being used to support those claims.

As I saw it, it was only in hindsight that observers could see the folly of Israeli leaders in 1973 and of American leaders at the Bay of Pigs, in Janis's case study. And it was only in hindsight that observers could see the wisdom of Israeli leaders in 1967's Six-Day War and of American leaders during the Cuban Missile Crisis, in another Janis case study. Knowing how things had turned out, they could confidently pick out the signals that these leaders had caught or missed – and come up with sure-fire ways to avoid such mistakes.

To me, this seemed wrong – and dangerous. In a phrase from my dissertation, "Those who perceive a surprise-free past are likely to face a surprise-full future." I recruited my friend, Ruth Beyth, to write something. Given that we were nobodies, we didn't even try to write a newspaper op-ed. Instead, we wrote a book review, of Victims

of Groupthink, which appeared in Megamot, the Israel educational journal, and, later, in English, in Policy Sciences.

About 10 years later, I met Janis, for the first (and only) time. Our book review was the only thing that he wanted to talk about. He thought that it overstated the case.

As I've thought about this, over the years, I've decided that he was right. Hindsight bias is very robust. It shows up in many places and is very hard to undo – unless one applies rigorous debiasing techniques, like those taught to historians. However, although hindsight bias pervades one's thinking, its effects on summary judgments are generally not enormous.

Averaging over many studies, in hindsight, an event that has happened looks like it seemed about 10% more likely than it actually seemed, in foresight. Sometimes, that change makes a big difference. For example, if it makes an event that seemed impossible look like it should have had a 10% chance of happening, that might support saying, "With a 10% chance, they should have acted." However, moving an event from a 30% chance to a 40% chance might not make much difference.

Indeed, one can even have hindsight bias bias, unfairly accusing critics of unfairly second guessing decisions. Moreover, at the extreme, if we act as though we can't be expected to predict anything, who needs us?

A couple of years after returning to the States, one of my doctoral advisors introduced me to someone who had played a central role in implementing the Agranath Commission's recommendation of setting up an independent analytical team. After about a year and a half, he had concluded that it couldn't work. His problems were not the social challenges of maintaining an independent stance, but methodological ones. As he put it, "Both teams have the same analytical approach, work with the same data, come from the same national culture, and survive similar institutional pressures. Why think that they're going to come up with anything usefully different?"

From a measurement perspective, one can learn something from the variability in the conclusions, when different experts apply the same method to the same data. At the extreme, if the "experts" don't agree at all, then we're all in the dark. Leaders need to know that, in order to take responsible gambles under conditions of high uncertainty. However, I saw his point.

What I didn't understand were his complaints about fundamental flaws in the intelligence community's methods. As he saw it, intelligence had two incompatible methodologies, each capturing some essential issues, while missing others. One method was the narrative analysis of others' intentions; the other method was quantitative analysis of others' logistic capabilities (that is, who could move what where and with what speed).

Although I have continued to work intermittently on security issues, sometimes with my friend, most of my work has focused on technological, health, and environmental risks. Indeed, for the last 19 years, half of my appointment has been in a school of engineering. Nonetheless, although the substantive content of those issues is very different from intelligence, the analytical issues are similar. Each domain faces the challenge of integrating complex, uncertain knowledge, distributed over diverse experts. Each domain has its own forms of narrative and quantitative analysis. Each runs the risks of groupthink, when forms of knowledge are restricted, by who gets to the table and what they are allowed to say.

Thus, one finds quantitative analyses of technologies' risks and benefits that are as elaborate as anything found in intelligence. The probabilistic risk analysis for a nuclear power plant might run millions of lines of computer code. The risk analysis for the ill-fated Shoreham, Long Island, nuclear plant was supported by an evacuation plan that reportedly ran 1500 pages.

One also finds narrative analyses of scenarios, like the best-case and worst-case scenarios that string together events that could occur, if everything broke right or wrong, in ways with some thematic unity. For nuclear power, the best case was seen in promises of "electricity too cheap to meter." A worst case was seen in "The China Syndrome."

Done well, both quantitative and narrative analyses can capture some essential aspects of complex situations. Moreover, both methods include procedures that can reduce groupthink. For example, modelers can elicit expert judgments in ways designed to overcome group pressure. Scenario development sessions can legitimate wide-ranging discussions.

However, both approaches can also exclude some kinds of information, frustrate external review and play poorly with one another. As a result, even exemplary group processes are inadequate, when groupthink is build into the analytical method and team.

Quantitative analyses tend to omit processes that are not readily quantified or familiar to members of the analytical team. As a result, human behavior tends to be poorly represented in quantitative analysis. It is easy to test the robustness of quantitative models to changes in assumptions about the issues that they address, nearly impossible to test for sensitivity to issues that they ignore.

In contrast, narrative analyses invite criticism. Anyone can raise problems with a scenario. However, it is often unclear what those problems mean for the scenario's validity. Does finding a single flawed assumption invalidate it? If so, how much of a change is required to accommodate it? If a scenario is accepted, how big a change in the world is needed to require its revision? For example, we have 15 National Planning Scenarios for coordinating homeland security activities. If we conduct a table-top exercise using one of those scenarios, can we tell how well we are prepared for it? Can we tell how well that prepares us for full the universe of hazards? It is easy enough for a

scenario to resonate to the group that created it, while being rejected by those who didn't share the experience.

I have often seen quantitative analysts gag at the imprecision of scenarios and qualitative analysts' eyes glaze over at the hyper-precision of models. Although sometimes not very gracious, these responses are probably the right ones. No one wants to lose cognitive control of the discussion around an important problem. And no one can change their ways of thinking very quickly. Those of us in the education business know how hard it is to give students an intuitive feeling for a good engineering model or historical narrative (or psychological experiment or mathematical proof or poem). It is a rare student – or scholar – who if fluent in more than one way of thinking.

Under stressful conditions, people are even less forgiving, regressing to comfortable modes of thought. For some, that will be "run the numbers for me, on what the results will be, if we follow this strategy." For others, that will be "walk me through how events will unfold, if we follow this strategy."

Decision makers want to get reports in the format that works best for them. Whatever form they get, they should have the assurance that its producers have done everything possible to incorporate other modes of analysis. The groupthink a modeling or scenario team may matter less than the groupthink embedded in those teams relying on just one mode of analysis.

In our own work, we have tried to integrate the two forms of analysis by combining *computable models* and *structured scenarios*. I'll describe it very briefly, to give a feeling what any integration entails.

A computable model is precise enough that one could, in principle, "run the numbers" and produce quantitative estimates. However, actual quantification is avoided, so as not to alienate qualitatively oriented users. Rather, the emphasis is on translating everything that matters to the experts into potentially measurable variables and relationships – the essential demands of computability. Doing so forces a level of precision that is often lacking in interacting groups – and in their communication with outsiders. It also creates a template for organizing information in a readily accessed way.

A structured scenario addresses all variables in the computable model, precisely enough to allow evaluating its coherence. However, such elaboration is avoided, so as not to alienate quantitatively oriented users, who will want to see the underlying structure. Rather, the emphasis is on ensuring that scenarios are consistent with the evidence incorporated in the computable model. Such scenarios allow people to express their intuitive views of the world in ways that can be translated into operational terms.

I apologize for offering a seemingly abstract approach in a talk without props. However, its application is actually quite straightforward. It doesn't require any special software or training – just a commitment to thinking in a disciplined way, perhaps

through multiple iterations, with a team incorporating all relevant forms of expertise. The last two entries in the Readings list below suggest how to do that.

Applications include evaluating warning labels for drugs, dietary supplements, and household chemicals; assessing the more esoteric risks and benefits of genomic health technologies; preparing for avian flu; understanding resistance to the anthrax and MMR vaccines; establishing transmission tariffs for deregulated electric power (in Ontario); and designing behaviorally realistic evacuation strategies after a terrorist attack.

Like their counterparts dealing with other complex, uncertain problems in other domains, intelligence analysts are inundated with information that they are periodically required to compress into the form demanded by decision makers. Like their counterparts, they need systems for managing knowledge that are compatible with both their ongoing task of making sense of their domain and their intermittent task of producing cogent summaries. Scenarios structured by computable models are one way to accomplish these linked tasks. The combination provides a knowledge management platform that can free analysts to deliberate in ways designed to avoid groupthink, confident that their conclusions can be given consistent model and scenario representation.

This past June, I was in Israel and had a chance for a long professional talk with my friend. In parallel, he had developed a remarkably similar approach, which he was applying mostly in defense-related projects. There are some differences. His models are less computable than I would like; ours are too ascetic for him. He typically maintains the models and scenarios for his clients, as circumstances evolve. I prefer to wrestle with the structure of a domain, then leave the product with others. However, I felt like our approaches were finally close enough for me to understand what was bothering him 30 years ago. It was, in my terms, the risk of groupthink in the choice of analytical method, as well as the risk in its application.

Readings
(available from Rosa Stipanovic, rl1y@andrew.cmu.edu or the author, baruch@cmu.edu)

Fischhoff, B. & Beyth, R. (1974). Failure has many fathers. Review of I. Janis, Victims of Groupthink., Policy Sciences, 1976, 7, 388-393.

Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 104, 288-299.

Fischhoff, B., Slovic, P. & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. Journal of Experimental Psychology: Human Perception and Performance, 3, 552-564.

Fischhoff, B. (1980). Clinical decision analysis. Operations Research, 28, 28-43.

Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S. L. & Keeney, R. L. (1981). Acceptable risk. New York: Cambridge University Press.

Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: State of the art to 1980. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases (pp. 306-334). New York: Cambridge University Press.

Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases (pp. 422-444). New York: Cambridge University Press.

Fischhoff, B. (1982). For those condemned to study the past: Reflections on historical judgment. In D. Kahneman, P. Slovic and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press.

Fischhoff, B. & Beyth-Marom, R. (1983). Hypothesis evaluation from a Bayesian perspective. Psychological Review, 90, 239-260.

Fischhoff, B. (1987). Do we want a better theory of deterrence? Journal of Social Issues, 43(4), 73-77.

Lanir, Z., Fischhoff, B., & Johnson, S. (1988). Military risk taking: C3I and the cognitive functions of boldness in war. Journal of Strategic Studies, 11(1), 96-114.

Fischhoff, B. (1989). Eliciting knowledge for analytical representation. IEEE Transactions on Systems, Man and Cybernetics, 13, 448-461.

Fischhoff, B. (1991). Nuclear decisions: Cognitive limits to the thinkable. In P. E. Tetlock, J. L. Husbands, R. Jervis, P. C. Stern, and C. Tilly (Eds.), Behavior, society, and nuclear war, (pp 110-192). New York: Oxford University Press.

Fischhoff, B. (1994). What forecasts (seem to) mean. International Journal of Forecasting, 10, 387-403.

Fischhoff, B. & Johnson, S. (1997). The possibility of distributed decision making. In Z. Shapira (Ed.), Organizational decision making (pp. 216-237). New York: Cambridge University Press.

Fischhoff, B. (2000). Scientific management of science? Policy Sciences, 33, 73-87.

Fischhoff, B. (2002). Heuristics and biases in application. In T. Gilovich, D. Griffin & D.Kahneman (Eds.), The psychology of judgment: Heuristics and biases. New York: Cambridge University Press.

Gregory, R., Fischhoff, B., Butte, G., & Thorne, S. (2003). A multi-channel stakeholder consultation process for energy deregulation. Energy Policy, 31, 1291-9.

Fischhoff, B. (2005). Decision research strategies. Health Psychology, 21(4), S9-16.

Fischhoff, B. (2005). Risk perception and communication. In D. Kamien (ed.), McGraw-Hill Handbook of Terrorism and Counter-terrorism (p. 463-492). New York: McGraw-Hill.

Fischhoff, B. (2006, May). Modeling: Visualizing your vulnerabilities. Harvard Business Review, 8-11.

Fischhoff, B., Bruine de Bruin, W., Guvenc, U., Caruso, D., & Brilliant, L. (2006). Analyzing disaster risks and plans: An avian flu example. Journal of Risk and Uncertainty. 33, 133-151.

BARUCH FISCHHOFF, Ph.D., is Howard Heinz University Professor, in the Department of Social and Decision Sciences and Department of Engineering and Public Policy at Carnegie Mellon University, where he heads the Decision Sciences major. A graduate of the Detroit Public Schools, he holds a BSc in mathematics and psychology from Wayne State University and an MA and PhD in psychology from the Hebrew University of Jerusalem. He is a member of the Institute of Medicine of the National Academy of Sciences, and has served on many NAS/NRC/IOM committees. He is a past President of the Society for Risk Analysis and recipient of its Distinguished Achievement Award. He has been President of the Society for Judgment and Decision Making. He is a Fellow of the American Psychological Society and the American Psychological Association, from which he received Early Career Awards for Distinguished Scientific Contribution to Psychology and for Contributions to Psychology in the Public Interest. He is a member of the Department of Homeland Security's Science and Technology Advisory Committee; the Environmental Protection Agency Scientific Advisory Board, where he chairs the Homeland Security Advisory Committee; the World Federation of Scientists Permanent Monitoring Panel on Terrorism; and the National Intelligence Council Global Expertise Reserve Program. He was a founding member of the Eugene Commission on the Rights of Women. His research includes risk communication, analysis and management; adolescent decision making; informed consent; security; and environmental protection. He has co-authored or edited four books, Acceptable Risk (1981), A Two-State Solution in the Middle East: Prospects and Possibilities (1993), Preference Elicitation (1999), and Risk Communication: The Mental Models Approach (2001).