

Appendix: Participant Papers

John Aherne

From Norman Bradburn's Workshop Plan:

"A primary goal of the workshop is to identify obstacles to effective use of existing research findings and tools, to suggest ways of overcoming these obstacles, and to identify crucial areas for future research."

Obstacles to *effective* use of existing research finding and tools:

- (1) Many of the findings do not address the needs of decision-makers, either because the research did not address the critical problems or because it is presented in a way that may be clear to the research community but is obscure to the decision-maker and his or her staff.

Communications between the holder of problems and those who can provide research answers must be improved. Funding agencies, such as NSF, may have to develop closer relations with mission agencies, the problem holders. It also may be necessary for funding to be provided so that researchers can spend time (a few weeks to a few months) in mission agencies to develop an understanding of the problems that need research. (This approach has been useful in environmental science research.)

Within a university, communication skills must be developed in the students taking risk courses so that, when they do complete projects, the work will be presented in a way clear to the recipients. Unfortunately, these skills also need to be developed in many professors.

- (2) Perceived and often real lack of objectivity. Risk analysis can be a powerful method to illuminate difficult choices. Situations where such choice must be made usually involve important funding or policy decisions, and, consequently, must be made under substantial pressure from those who already have decided what is the correct choice. Conflict-of-interest charges can be used to discredit analyses that differ from such "correct choices". Unfortunately, enthusiastic professors also may "know" the correct answer and convince students that objectivity need not be an overriding practice. Interdisciplinary teams are less likely to succumb to prejudice and can provide mentoring to students on how to do objective analysis of difficult problems. Providing funding by objective agencies, e.g., NSF, rather than industry or mission agencies also provides greater independence and less perceived or real pressures for specific results.

- (3) Inability or unwillingness to understand the complexity of decision analyses. Most decision-makers rely on staff members to develop answers to problems. In all three branches of the federal government, senior people are advised by staff that usually are only a few years out of graduate or professional school. Many of these staff have not had education in the disciplines used in risk analysis. While quite bright, they need help to understand how to work through sophisticated analyses.

The suggestions for (1) can help here. More beneficial would be to incorporate courses in applying risk analyses within the graduate and professional programs, particularly in those schools from which come the majority of interns and new staff

Crucial areas for future research:

- (1) What are the synergistic effects of exposure to several pathogens or to a combination of chemical and biological or chemical and radioactive hazards?
- (2) What are the effects on emergency response systems when compounded by large-scale panic in addition to physical challenges to humans?
- (3) How can economic analysis into hazard analysis be combined to develop more credible cost-effectiveness studies?
- (4) How can uncertainties and probabilities be presented in ways that lead to understanding by the general public and decision-makers?
- (5) Where in undergraduate and graduate departments should interdisciplinary courses on risk analysis and decision-making be housed?

Another important issue:

From Bob O'Connor's 2 May email:

“Ensure that the best researchers have appropriate incentives to participate”

The “best researchers” do not have difficulty getting funding for themselves or for their students. To convince those researchers who are not now working on the risk-related problems to shift their efforts will require

- Interesting and important problems.
- Sufficient funding for themselves and their students.
- Confidence that the funding will continue to be available for several years (or else it would be irresponsible for a professor to ask students to do their work in the area).

David A. Asch

Individual vs. Population Risk

Most of the examples put forth as problems in risk analysis and decision making are risks that populations face—or at least that make little distinction among the sometimes different risks that individuals within those populations face. Shortly after the bombing in Oklahoma City, the political cartoonist Tony Auth produced a cartoon showing a map of the United States with dots across it representing various towns and cities. The way Auth drew the map, all of the cities had same name: “Oklahoma City.” The cartoon reminded us that the attack could have happened anywhere, that we are all vulnerable, and also that we all shared in the feelings of loss. Similar feelings arose after the September 11 attacks. The direct victims were on airplanes, at the World Trade Center, and at the Pentagon—but the attack was perceived as against America, and the persisting sense of vulnerability has been national.

There are social, political, and practical reasons why it is easier to think of risks at the level of populations rather than at the level of individuals. Population risks are relatively easier to calculate. They would seem more amenable to efficient “policy solutions.” They also depersonalize risk and its mitigation in a way that may make the acceptance, understanding, and potential reduction of those risks more feasible.

In health contexts, risks are seen sometimes from a population perspective and sometimes from an individual perspective and, if anything, the individual perspective seems to be increasing. At the population level, dental caries were seen as risks to all, and widespread (but not universal) water fluoridation was introduced. As evidence mounted that high blood pressure was silent, prevalent, dangerous, and treatable, mass screenings were instituted that affect all members of the population. All women within selected age groups are encouraged to undergo regular mammography.

The individual perspective hasn’t been seen as explicitly competitive with the population perspective, but it has increased for several reasons. One reason is that the methodologic techniques of epidemiology employed to understand population health risk have also proved attractive for investigating individual risk factors. Another reason has been the political ascendancy of the medical model emphasizing individual medical interventions as the path toward health. This model was often supported by the power and infrastructure of US schools of medicine and contrasted somewhat with the reduced finances and influences of schools of public health. Another reason has been the increased attention on human genetics and with it the promise of highly individualized health risk identification and health risk mitigation.

For example, cardiovascular risk reduction is now considerably more individualized than many other disease preventing interventions that are more population based. An individual patient’s “target” for cholesterol (or subfractions, like LDL cholesterol), has been positioned as a function of age, sex, family history, other lipid fractions, co-morbidities, smoking history. The one-size-fits-all model (that, for example, everyone should take low dose aspirin) has been displaced by a model that, if not based on genetic testing, is at least highly individualized.

Although the vast majority of breast cancer is not attributable to the *BRCA 1* and 2 mutations, breast cancer risk is also increasingly individualized based on mutation status, age, family history, and personal reproductive history—so much so that not only are chemoprophylaxis and surgical prophylaxis decisions based on these factors, but authorities are reconsidering the approach to screening, moving it from its current population stance, to one that is more individualized.

The logical extension of this individualization is that the risk each person faces is separately calculable, that each individual has his own targets for risk reduction, and that there are no real public health risks or public health interventions—only collections of individual risks and interventions. These approaches present the following issues:

Information integration. The problems of information integration are substantial enough when population central tendencies are used to make policy recommendations. Progressively more individualized decision making progressively increases these cognitive burdens for both patients and physicians. Innumeracy among both patients and physicians further challenges the practical implementation of individualized (or population-based) risk reduction. Automated decision aids could be part of the solution, and some exist to help individuals negotiate their breast cancer risk, for example, but the development, use, and regulation of these aids could benefit from the kind of experience convened for the workshop. Absent these efforts, substantial paternalism for risk management decisions might be the only other seemingly practical approach, and yet are largely intolerable in the US health care setting. (Along with putting fluoride in the water, perhaps we should also add folic acid, aspirin, beta blockers, statins, and Prozac.)

Crises of confidence. Trust in institutions (e.g., the medical or nursing community) plays such an important role in both individual and public health risk management. Lack of public trust in health settings would be paralyzing given the high degree of individual cooperation required to effect risk mitigation. But confidence in these systems is extremely fragile: as of today, many are suspicious of *all* accountants, CFO/CEOs, and Catholic priests. In the health setting, the recent debate about the actual value of mammography screening has aired in the lay press as well as in scientific journals. Will that debate reduce mammography uptake? (Maybe it should.) More individualized risk assessment might further increase perceived ambiguity (if second order uncertainty increases in the setting of risk groups). Effects might be overgeneralized to all health interventions.

Time demands. Although in some ways the distinction between individual and population risk assessment and management may appear to be a false dichotomy (after all, populations are collections of individuals), individualized decisions create huge aggregate time demands. The current challenge in thinking about a policy for population or subpopulation smallpox vaccination is nothing compared to how that thinking will need to play out over and over again as it is *reconsidered* at the individual level. At least in the US, individuals are not accustomed to giving away authority for health decisions that are implemented personally. Even if individual patients ultimately accept default or modal recommendations, they feel entitled to individualized discussions that contextualize the decision within their particular situations. Even if those discussions could be limited to 30 seconds, five 30-second discussions of this type, discussing five different interventions, would overwhelm a primary care encounter. And the aggregate burden of these activities across patients would be enormous.

Intolerance for individualized risk. Despite some pressures to individualize risk assessment and its management in health settings, there is also increasing intolerance of these activities because they compete with social and political forces that are currently very powerful. Individualization of risk implies progressively segmenting the population into determinate groups based on observable characteristics. Not only does this individualization make it easier for notions of blame and stigma to be associated with risk (and thus perhaps harder to increase acceptance of risk or even reveal it's source), but it also evokes socially unacceptable traditions. Racial profiling in the setting of local law enforcement has been roundly criticized. Racial profiling in the setting of anti-terrorism initiatives has also been criticized, but probably less so. In health settings, many react with horror at the idea of limiting screening for sickle cell disease to the African American population. Such specification is considered “targeting”—a term sufficiently charged and loaded that it might overwhelm competing arguments from efficiency (that despite imperfect determinations of race and increasing genetic mixing, sickle cell mutations remain relatively rare in non-African Americans). It is hard to be Bayesian and politically correct at the same time.

What problem areas would be appropriate for consideration in a new integrated research program on risk analysis and decision making?

In summary, I thought many of the ideas presented in the other background papers were much better and more general than mine, but I came in late and this is the best I could do without repeating the contributions of others. Nevertheless, I think the movement from *population*-based risk assessment, risk perception, decision-making approaches and risk management strategies toward *individual* approaches to all of these is particularly relevant to health settings and may be more generally relevant as well.

For the above problem areas, what are the most significant research challenges that we should consider?

I'd have to pick information integration as the most formidable challenge in terms of implementation—largely because untrained individuals have such a hard time manipulating and incorporating risk information (as do trained individuals). The main research question in this area I find of interest is contrasting appropriate approaches for group decision making versus individual decision making.

How can an NSF initiative most effectively promote future research in risk analysis and decision making that brings together the natural and social sciences?

I think there has been the view that NSF should steer clear of health-related issues because NIH and other health-related agencies cover that area. But with the exception of some programs at NCI, very little at NIH, it seems to me, can really address what is effectively the basic science of risk analysis and decision making. At the same time, so much appropriate substrate exists in the health domains—substrate at both the population health and the individual health levels. And NIH and other health-related mechanisms have created a huge infrastructure of scientists and protocols, clinical trials, epidemiologic studies, health services research assessments. Any of these could have basic work on risk and decision making incorporated at marginal cost and—importantly— within increasingly naturalized settings. These naturalized settings would add credibility and impact to the field and might bring in new investigators.

Vicki Bier

- What problem areas would be appropriate for consideration in a new integrated research program on risk analysis and decision-making?

I certainly agree with the focus on homeland security and natural disasters. I also believe that research on catastrophes is important, regardless of whether they are due to natural disasters. One area not specifically called out in the workshop plan, but which I believe merits further work, is the field of environmental risk management. Finally, further work on the effects of income and class on health would also be beneficial.

- For the above problem areas, what are the most significant research challenges that we should consider?
 1. Specific research questions to be addressed in the area of **homeland security** might include the following:
 - A. People may have a legitimate desire to be informed about protective measures being undertaken to enhance homeland security. However, making information about protective measures known will often reduce their effectiveness. An example is the installation of sterilization equipment at post offices, to protect against anthrax delivery through the mail. Such equipment might be effective if it could be installed secretly. However, if the installation of the equipment is public knowledge, it will just deflect the threat to some other delivery mechanism (e.g., Federal Express). How can the strategic need for secrecy be balanced against the public right to information? Are there risk communication strategies that can provide useful information to the public, but not give away strategically important information?
 - B. As pointed out in the workshop plan, better security of one target may merely shift terrorists' efforts to another target. This makes it important to distinguish between technologies that provide a public benefit, and those that provide primarily private benefits. For example, if defending a major corporate headquarters simply results in efforts being shifted to the headquarters of a different corporation, then little if any public benefit will have been achieved, despite the expenditure of possibly large amounts of money. By contrast, investments that might yield a substantial public benefit would include technologies sufficiently inexpensive for widespread use, technologies that are mobile (e.g., enhanced first response capabilities that can be mobilized after an attack), and technologies that provide benefits above and beyond homeland security (e.g., enhanced identification of disease clusters, regardless of cause). Research to identify projects likely to provide substantial public benefit could be an important building block in establishing national priorities for homeland security research in general.

2. The low-probability, high consequence nature of **catastrophic events** makes them a good test bed for risk assessment and management methods. Bier et al. ("A Survey of Approaches for Assessing and Managing the Risk of Extremes," *Risk Analysis*, Vol. 19, No. 1, 1999, pp. 83-94) provides directions for future work in this area. Specific research questions to be addressed in this area might include the following:
 - A. Further research on which approaches to decomposition yield the best results would be useful. There may also be some classes of problems for which useful decomposition or modeling approaches have not yet been identified, or for which severe computational problems exist. Further work in such areas may yield feasible and beneficial modeling approaches.
 - B. The risk analysis community has paid little attention to recent developments in Bayesian statistics, such as the work on reference priors and robust Bayesian analysis. Much work has been done in these areas, from which the risk analysis community might benefit. The field of econometrics in particular has pushed the development of rigorous Bayesian methods to a level of sophistication not usually seen in risk analysis. Further work extending such methods to address issues of concern in risk analysis (such as the small amounts of data that are typically available on extreme events) might be desirable.
3. In the area of **environmental risk management**, I would focus on problems with irreversibility, non-fungible effects, large uncertainties, and long time horizons. Examples might include the following:
 - A. Risk management for global warming. Arguably, due to the long time lags involved, if global warming were real, then in order to be effective, actions to mitigate its extent would need to be adopted before we could even be sure that the phenomenon was real. Analysis similar to that undertaken by Carnegie Mellon and other groups could help clarify the circumstances under which immediate action should be taken, and those in which further research is preferred.
 - B. Species extinction (including the possibility of human extinction due to nuclear winter) is an example of a phenomenon where irreversible and non-fungible effects are clearly important. Such problems pose both significant ethical challenges, and also technical challenges to existing decision-making approaches such as discounting. For example, Mannix has pointed out, in a review of *Discounting and Intergenerational Equity* by Portney and Weyant ("How Much Do We Care about the Deep Future?" *Regulation*, Vol. 23, No. 2, pp. 55-57), that "There is no market for planets." Therefore, the benefits accrued over time from investing in productivity enhancement rather than environmental preservation cannot necessarily compensate for severe environmental devastation. In fact, the fundamental basis for discounting is the idea of productivity growth, which could cease to be a reasonable assumption if the carrying capacity of the earth is seriously diminished.

4. There is a rapidly expanding literature on the unintended consequences of imposing regulations, by Keeney, Graham, Viscusi, and others. So far, this literature has tended to indicate that the economic costs of regulation may increase risk by reducing household income. However, I am concerned that these analyses may not take adequate account of confounding factors.
 - A. In particular, I think there is a need for further work to more effectively control for factors such as class and values in such analyses. For example, even if poor people as a rule pay less attention to health preservation than middle class people, is this primarily because of income differences, or largely due to class and value differences? If middle class people become poor, do their expenditures and values change to reflect their new income level, or are middle class values such as health preservation a legacy that survives changes in income levels?
 - B. More work is also needed on the role of regulation in reducing the cost of protective measures, by expanding the market for such measures, thereby encouraging innovation. Therefore, some regulations could enhance social welfare in the long run, if the reduced cost of protective measures makes them more affordable, even if the short-run effects of the regulation would appear to be undesirable.
 - C. I believe there is also a role for further research on the distributional effects of environmental cleanup efforts. In particular, under what circumstances do such efforts merely transfer risk from the middle class communities that are being cleaned up, to increased (largely occupational) risks for working class communities?
- How can an NSF initiative most effectively promote future research in risk analysis and decision-making that brings together the natural and social sciences?

I believe that many of the problems described here would be well addressed by large, cross-disciplinary, multi-institutional teams. However, I am not convinced that the existing mechanisms at the National Science Foundation for funding collaborations—Science and Technology Centers (STC) and Engineering Research Centers (ERC)—are well suited to this purpose. In particular, the critical contributions of the social sciences to these problems are likely to be under-appreciated by the review panels that make decisions regarding the existing STC and ERC programs.

Therefore, I would encourage establishment of a specific grant program for integrated research in risk analysis and decision-making. Such research could be jointly funded with various mission agencies (such as the Environmental Protection Agency), on a case-by-case basis as appropriate. The grant amounts might not need to be quite as large as for the existing STC and ERC programs, since much of the work would likely not require large amounts of laboratory equipment. However, they should be sizable enough to facilitate cross-disciplinary and multi-institutional research, rather than smaller grants for one or a few investigators.

Robin Cantor

I believe that many other participants in the workshop are far more qualified than I am to speak to the first and second questions posed by the workshop description. My contribution will therefore largely focus on the third question, “How can an NSF initiative most effectively promote future research in risk analysis and decision making that brings together the natural and social science?”

I believe a major obstacle to promoting future interdisciplinary research in risk and decision making comes from *within* the research community. In my view, these research areas are undervalued by the university and scholarly communities. The Society for Risk Analysis has considered this problem and circulated a statement regarding the funding of a number of university-based risk centers. The SRA proposal focuses on using centers to make a concerted national investment in the long-term institutional infrastructure of risk analysis. These investments are intended to elevate the academic standing of risk analysis through improvements in the educational and scholarly components of risk analysis within the university structure. At the same time, these investments are expected to expand our national capabilities to respond effectively to risk problems.

Implementing a new investment in risk centers, however, is a formidable task. The program announcement must signal the key objectives to be addressed by a center proposal, which are likely to differ substantially from the criteria for PI research proposals. The selection process is likely to be very different as well, with much more emphasis on proposed methods to achieve educational and scholarship functions, conduct outreach, and enhance exchange of information and people across institutions. I hope that the workshop participants will spend some time discussing the many difficult questions that arise in the design of an appropriate request for center proposals and in the creation of an appropriate review mechanism to judge center proposals.

The rest of my statement is an abbreviated version of the SRA Council’s statement about the needs and benefits of the center proposal.

ABBREVIATED STATEMENT OF THE COUNCIL OF THE SOCIETY FOR RISK ANALYSIS

Overview

In this statement, the Society for Risk Analysis (through its elected governing council) recommends that President Bush and the United States Congress establish a national, university-based program of interdisciplinary training and research in risk analysis. The SRA Council believes that expansion of university-based activities in risk analysis and risk management is crucial to improving the decision processes and the strategies for reducing losses and providing funds for recovery from future extreme events, whether they are malevolent or random occurrences. These catastrophic risks can be related to environmental and ecological, technological, food safety, medical or natural hazards and/or the risks associated with terrorism to the Homeland. The end result will be the potential improvements in safety, efficiency, and equity regarding these and related issues.

A concerted and focused investment in the educational and research infrastructure supporting risk analysis will produce the foundation for better public and private responses to a wide range of high priority concerns including but not limited to infrastructure protection, drug development and marketing, hazardous waste management, food safety, airline security and safety, pesticide use and regulation, energy security.

The Council of the Society for Risk Analysis is concerned that the demand for qualified professionals and scientists will soon outstrip the number of well-trained people in this field. Professionals who understand and can integrate the various components of risk analysis for a wide range of problems will be increasingly in short supply. Such shortfalls may ultimately constrain growing reliance on risk analysis approaches. Because universities tend to be dominated by scholars who have a traditional disciplinary orientation, the Council believes a concerted national investment in long-term risk capabilities is warranted. This national response can address directly and specifically the need for more and better educational opportunities in risk analysis and the critical need for linking science and engineering with social science and policy-related research.

The Educational and Research Challenge

The SRA Council has identified four levels of educational needs in risk analysis within universities:

- 1) Undergraduates majoring in science, mathematics, engineering, economics, and other social sciences need access to a high-quality, interdisciplinary course on the subject of risk as well as to modules on risk that are integrated into current required courses. Such courses and modules would teach undergraduates to identify, describe, and evaluate risks. In turn, this would help students better understand the implications of their own personal decisions, and make more effective professional choices.
- 2) Graduate students in numerous fields (particularly medicine, public health, statistics, law, environmental studies, toxicology, pharmacology, ecology, engineering, economics, and the social sciences) need to integrate the principles and methods of risk analysis into their required curricula.
- 3) Mid-career professionals in the public and private sectors, including government, business, the mass media, and advocacy organizations, need access to short yet intensive educational opportunities that can upgrade their understanding of risk-related methods and issues.
- 4) Doctoral-level training in each of the three facets (assessment, management, and communication) of risk analysis also needs to be provided to a small but growing number of leaders who can advance research in the field and develop the next generation of theory, curricula, and applied approaches.

Although some universities have developed useful programs in these areas, the Council believes that at present no university in the United States is capable of providing the high level of education in risk analysis outlined above, stimulating the requisite interdisciplinary integration

on risk issues, and creating value-added opportunities for multi-university or regional perspectives. Many current efforts, although promising, do not embrace the broad, interdisciplinary approach that is critical to successful education in this field. While some of the existing programs are interdisciplinary and involve talented, productive faculty, they are understaffed in terms of the level of resources needed to mount these programs at the necessary scale.

A related concern is the capacity to conduct research to continually improve on the concepts and methods of risk analysis. While a number of productive scholars are involved with the existing university programs in risk analysis, virtually no program has the staff to mount anything close to the scope of new research projects that should be conducted to address the range of needs and opportunities in this field. All the existing risk analysis programs are small compared to other disciplinary programs in universities. As a result, initiatives that link scholars across universities into research project teams and support the activities of a new crop of researchers in post-doctoral and graduate programs have great potential to build risk capabilities and resources to meet both short and long-term needs.

A National Program for Training and Research in Risk Analysis

In order to address these needs, the Council of the Society for Risk Analysis recommends that the Administration and Congress create a national program for training and building research capacity in risk analysis. At the outset, this would mean supporting several (up to ten) high-level university-based centers in risk analysis in the United States. A substantial number of centers are needed to satisfy regional needs throughout the country. The SRA Council recommends that each Risk Analysis Center be funded at approximately \$1 million per year for a five-year start-up period. At the end of the start-up phase, the program would be reevaluated. The SRA Council believes that this level of support would be adequate to stimulate significant innovations within participating universities, but it is not excessive in light of the current amount of faculty expertise now present in universities throughout the country.

Where appropriate, these Centers might attempt to fill broad regional needs and build regional expertise by serving a constituency wider than a single university. Multi-university teams could be assembled to seek Center funding. One result could be curricula (for any of the four educational contexts outlined above) shared among several universities that would allow students to take courses via distance-learning technologies. Thus, one Risk Analysis Center could serve a wider audience by providing students at other universities, and professionals at remote sites, with opportunities to participate in courses of study.

The university-based centers for risk analysis should be designated and funded through a competitive, peer-review process that is typical of the National Science Foundation. In evaluating proposals to launch such centers, the Council recommends that appointed review panels seek the following features in successful proposals:

- evidence of outstanding senior faculty leadership in risk analysis, or substantial faculty-level expertise in at least one of the three sub-fields of risk analysis (assessment, management, and communication), with a five-year plan to incorporate the missing fields into the Center's core faculty;

- clear commitment to an interdisciplinary approach to risk analysis education;
- evidence of how currently operating curriculum offerings in at least one of the four levels cited above (undergraduate education, graduate/professional education, mid-career education, PhD education) are subjected to serious evaluation and refinement, and a five-year plan to integrate risk-based educational offerings into the university's educational products;
- evidence of ongoing interdisciplinary research projects on important risk-related topics with substantial involvement by diverse faculty and students;
- plans to establish an advisory committee of practitioners from the public and private sectors that can bring a sense of real-world problems into the university environment;
- evidence of significant commitment to risk analysis from the applicant university as well as significant support for risk-related activities (training or research) from interested sources in the public or private sectors; and
- a strategy for and a commitment to making the proposal extend beyond a single university to address such factors as regional needs and strengthening regional and national networks, information resources, and intellectual capabilities... Examples would be a commitment to share and transfer risk-based educational and training materials through multi-university partnerships, or involvement of faculty and researchers from several universities or research institutions in the proposal.

The funds provided through this program could be used for three major purposes:

(i) developing and implementing new educational programs in risk analysis. Specific activities or budget items could include: creating advisory panels or conducting surveys to help guide design of new educational programs, curricula development (which is quite costly when distance education is one objective), efforts to better integrate research findings into curricula, support for center administration (including efforts to link faculty from various universities), support for mounting the new educational programs, junior faculty development, and mid-career retraining for faculty.

(ii) building research capacity in risk analysis. Specific activities or budget items could include: support for post-doctoral fellows, partial support for adjunct researchers, support for graduate student research projects, support for efforts to link researchers from various institutions into research teams, support for proposal development, seminar series, strategic planning for research programs, and key physical infrastructure items.

(iii) providing graduate student support in risk analysis. Specific activities or budget items could include: partial support for tuition and living expenses, support for thesis research projects (noted above), support for creating and implementing professional internships, and support to attend conferences.

The ultimate product of the proposed centers will be a new generation of faculty, professional practitioners, graduates, and students with the knowledge and skills to apply risk analysis in ways that can improve our ability to responsibly address national priorities.

Carl Cranor

Background Issues:

1. Normative Issues in Assessing Decisions of Risks and Uncertainties

Paul Slovic has noted that an important element in assessing risks are intuitive approaches to the issues. There are also normative considerations, typically that would be discussed by philosophers, in addressing risks and the acceptability of risks. These include such things as the source of the risk (naturally v. humanly caused), the relation of risks and benefits to the risk bearer's life (central v. peripheral; whether the risk bearer is a substantial beneficiary or not), the degree of voluntariness of risk-exposure, including how transparent and palpable the risks are, the personal avoidability of risky outcomes, and the costliness of taking the options; the extent to which a risk-bearer participated or not in decisions that left him or her exposed to risks; and the extent to which the risk-bearer can exercise continuing control over the risks and their materializing; and the reliability and legitimacy of a trustee for the public, such as a regulatory agency, in assessing and controlling risks.

All of these considerations are important in being able to justify risks to risk-bearers, something from a moral or social point of view to which we must attend. In addition, there are different global (philosophers would say "moral") principles for adjudicating the social or moral acceptability of risks in a community. Much of the social acceptability of risks in contemporary U.S. discussions is dominated by one moral paradigm, utilitarianism or its offshoots, risk-cost-benefit analysis. There are other views that receive much less attention that can provide insights: justice theories and other views concerned with the distribution of benefits, harms and risks in the community, and these in turn have variations. Some variations might be concerned with maximizing benefits to those who are worst off in the community, with persons having equal standing to participate in decisions that affect their lives, with the antecedent well-being of those exposed, with the susceptibilities of particularly vulnerable individuals or with heavily impacted sub-communities, or with other distributive considerations.

The above are merely examples, but considerations that frequently are not considered in discussions of the acceptability of risks. These considerations are pertinent, not only to discussions of the acceptability of risks, but also can help inform risk communication. Many of the conclusions to which I would come utilizing many of these considerations would agree in large measure with the results of research done by Paul Slovic and Baruch Fischhoff, but would have a different foundation. Moreover, I also tend to agree with a number of the remarks Paul has forwarded for this workshop, but again they would have a different rationale.

2. Risks, uncertainties and ignorance

A second background point, echoing those of some others, is that not all situations concerned with risky outcomes can be analyzed with risk analysis. Some involve great uncertainty; others considerable ignorance, so that one cannot easily assign probabilities to outcomes and may not even know the range of possible outcomes, e.g., the ecological effects of humanly created substances into the environment.

Researchers and decision makers need to be more candid when such circumstances obtain and perhaps NSF needs to encourage researchers to develop methods and procedures for addressing them. How much research has been done on addressing circumstances in which considerable ignorance obtains, I do not know (its not my area of research), but the discussion largely appears to have been dominated by “risk analysis” in which both a reasonable range of outcomes could be defined and one could provide reasonable estimates of risks for many or most of the outcomes. As the range of outcomes becomes more vague or completely unknown and it is less plausible to assign probabilities to any outcomes that are known, the problems for social decisions become much more difficult. In such cases, decision makers need guidance on what to do. What normative advice does one provide them and what does the social science research on risk perception and risk communication suggest the larger public thinks about such circumstances?

With 1) and 2) as background I have the following further observations on some of the issues posed for the workshop.

- I agree with the comments under research questions that we need guidance for addressing decisions under great uncertainty or ignorance.
- I think we should be careful about saying that that the public has been “caused” to have different assessments of risk (that could be) or acceptability of risks. This suggests that the public is in some way mistaken or irrational (they might be concerning probabilities of outcomes, but they might be addressing different issues such as the acceptability of the risks and outcomes). Moreover, there may be quite good reasons for the public to have quite a different assessment of the acceptability of risks than professional risk assessors (I agree with Paul on this). In such circumstances one should explore more carefully the basis of disagreement to see whether and what plausibility to the views the public might have.
- Similarly, when there is a difference between the public’s perception of risk (does this mean the “acceptability of risks”?) and scientific assessment of (the acceptability of?) risks, one must look closer so as not to dismiss the public’s view.
- On the issue of transboundary risks or risks associated with trade, it is perhaps useful to consider alternative approaches to addressing them than with an insistence on strict risk assessment methods. Given 1) above a public in one country may simply not want to live in a world with certain low level risks, even if another country does. This is a perfectly

reasonable choice in a democracy. How does one address such issues? I have no answer, but I would urge that one not simply dismiss one public as being irrational or misguided vis-à-vis another public; they may simply be making different normative decisions about how to live their lives.

- On equity and distributional issues see 1) above.

3. The three questions:

- At this point I have no particular suggestions, if this means particular or generic risky situations, about what should be addressed.
- Some of the research challenges include addressing the issues suggested in 1) and 2) above.
- I would urge NSF not only to incorporate natural and social science, but also to include those who, like philosophers, can address some of the normative issues raised in 1) as well. NSF could make it a condition of funding that research teams should include a range of natural and social science experts as well as some who are expert in the normative issues as well.

Susan L. Cutter

Risk Assessment or Vulnerability Science?

The events of September 11th highlighted in very graphic terms the necessity for understanding vulnerability to extreme events. Currently, our knowledge base is fragmented and insufficient to advance our understanding of many key issues in risk or hazards assessment. For example, the research communities are often too focused on local social dynamics and impacts (post-event qualitative case studies), or solely engaged in the development of physically based models of the threat (where we can determine it). What we lack is the conceptual and analytical ability to integrate across these domains. Therefore, we need a new approach or framework for gathering knowledge to guide and advance risk, hazards and disasters research.

Vulnerability science (the potential for loss in human, natural, or built systems) is one such approach. Depending on the perspective (or discipline) vulnerability is the potential exposure (who or what is at risk), or an indicator of the coping capacity of the environment, individuals and society to suffer harm including their resistance, resilience, and recovery from risk or hazard events. Vulnerability science is a broader concept than risk analysis, yet builds on the existing interdisciplinary and integrated tradition of risk, hazards and disaster research, incorporating both qualitative and quantitative approaches, local to global geography, historic to future temporal domains, and best practices. It also adds more technological sophistication and analytical capabilities to our current research methods, especially in the realm of the spatial sciences (e.g. GIS, spatial decision support systems). A research agenda for vulnerability science was recently proposed and could serve as a useful starting point for discussion (S.L. Cutter 2001. "A Research Agenda for Vulnerability Science and Environmental Hazards," IHDP Update, Newsletter of the International Human Dimensions Programme on Global Environmental Change 2/01: 8-9).

Current Constraints on Risk and Hazards Assessment

1. There is no national inventory of natural hazard events and losses although pieces of such a database exist throughout the federal agencies.

The lack of such a reference baseline for documenting hazards events and data on losses (economic, human casualties) at a sub-national level, precludes the development of a simple all-hazards risk or vulnerability assessment for particular places. Further, there is inconsistent measurement of economic losses and costs, where distinctions between direct and indirect, tangible and intangible, are not made in any coherent fashion. Data on non-natural events and losses (including terrorism and the suite of technologically-driven hazard events) are even worse. Therefore, we need a systematic accounting of hazard losses by specific hazard source by location for the nation in order to document the past history of losses and use this as an analog to develop a predictive model of potential future losses (S.L. Cutter (ed.) 2001. *American Hazardscapes: The Regionalization of Hazards and Disasters*, Washington D.C.: Joseph Henry Press).

2. We need to know who is at risk, where the risk is distributed, how often the risk manifests itself, and the cumulative effects of multiple risks and hazards on people and places.

At present we lack sophisticated models of biophysical risk exposure and societal risk exposure that are linked and integrated with one another. There are three kinds of risk assessment models: risk estimation (the probability of occurrence and its spatial extent); those that examine exposure assessment (physical environment, building stock, financial) and those that examine the impact of an event in the form of loss estimation (normally direct economic losses and human casualties). Most risk assessment models are hazard specific (e.g. airborne contaminants, floods, earthquakes) and geographically or temporally restricted. Very few take into consideration multiple stressors and impacts at different scales (ecosystem to global, neighborhood to nation). Various intellectual communities are involved in the development of these models—natural sciences, social sciences, engineering sciences, and as a consequence there is little overlap or understanding of each (NOAA 2001. Proceedings of the Forum on Risk Management and Assessments of Natural Hazards. Washington D.C.: Office of the Federal Coordinator for Meteorological Services and Supporting Research). FEMA has developed HAZUS (Hazards US) as a tool to forecast economic impacts from seismic hazards (wind and flooding modules are in development), which integrates risk estimation, exposure assessment, and loss estimation parameters within a Geographical Information System, but again this is limited to a single hazard type.

3. We do not know how current risk and hazards assessment methods and practices contribute to the relocation of risk and vulnerability (either geographically or into the future).

Some of the methods used to assess risks and hazards may include social or spatial biases. For example, depending on which toxicity indicator is used, the delineation of the potential exposures to the community may be highly variable. This is especially important in environmental equity considerations. We need a more robust set of techniques and tools that allow us to examine these locational inequities and geographically represent risk, hazards, and vulnerability. These equity concerns are inherently geographical, yet there are limited research opportunities to experiment with different forms and modes of visualization (e.g. a map; an animation, virtual reality), which can enhance both the public's and decision maker's understanding of the differential impacts of risks and their management. We also need better theoretical and conceptual understanding of the process and impacts of risk transference (e.g. cleaning up Rocky Flats and moving the plutonium to the Savannah River Site in South Carolina, or the transboundary shipments of hazardous materials from developed to developing nations).

4. We need new methods for representing uncertainty in our forecast information and different ways to visualize and represent surprise in some of our conceptual models of risk, hazards, and vulnerability.

This new knowledge must be developed through the integration of the various research communities (risk sciences, social sciences, natural sciences, and engineering sciences) and incorporated into the development and training of the next generation of researchers. At present, there is very little interaction among the risk, disasters, and hazards research communities. Stimulating the cross-fertilization of ideas and perspectives should be a priority.

Suggestions

- 1. Establish a national clearinghouse for hazard event and loss data that is partially supported by NSF, an effort that would provide the baseline data for all-hazards assessments by geographic locale (county) for the nation.**
2. Initiate a research focus within NSF on integrated models for risk estimation, exposure assessment, and loss estimation in order to improve the science in support of risk management policies.
3. Create a competition for Centers of Excellence in Risk Analysis and Decision Making (5 year funding maximum) in order to stimulate interdisciplinary research (and the training of students) to address some of these critical deficiencies.

Baruch Fischhoff

Integrating Risk Analysis and Risk Communication

Many bodies in this country and elsewhere have called for integrating risk analysis and risk communication, as a necessary condition for ensuring the credibility and validity of each activity. Doing so poses cognitive, organizational, and analytical challenges. It means consulting citizens early enough in the analytical process and with sufficient precision for them to shape its terms. It may require finding analytical representation for hard-to-model concerns (e.g., psychological impacts, equity, human error). It means expanding the envelope of public participation, while realistically assessing its success and judiciously describing public competence.

The political acceptability of the analytical enterprise depends, in part, on its ability to find suitable roles for the public – both as sources of input and recipients of output. The current skepticism about analysis (e.g., as seen in some calls for precautionary principles and some opposition to regulatory reform) has multiple sources. These include (in no particular order) innumeracy, scientific illiteracy, poor expert communication, specific analyses producing unwanted results, lack of resources for critiquing analyses, and the perception of overly narrow analyses and captive regulatory processes. The supporters of analysis have their own complex strengths, weaknesses, and motives.

Analysts inevitably take some positions on these issues (e.g., by whom they consult, by how they define terms – like “risk,” by when they communicate their results). The research challenge is expanding the range of possibilities for analytical processes. The current calls for expanded public involvement represent a potentially important social experiment, shaping the roles of citizens and technical specialists. The success of this experiment will depend, in part, on technical execution. Amateurish efforts could sour the participants on the experiment and on one another.

The plans for this research should include strengthening links between the risk world and the traditional academic disciplines. In terms that Alan Baddeley once suggested, there is the opportunity here for applied basic psychology (testing the generality of basic research results) and basic applied psychology (extracting general phenomena from applications). On the one hand, risk buffs need all the help that they can get – as well as pressure against spinning off into their own self-contained world. On the other hand, the core disciplines should benefit from having the new problems – and from the pressure to integrate applied and basic research. Society should benefit from having a stronger connection between risk-related activities and peer-reviewed research. Since September 11, there seems to be somewhat greater interest in these issues, regarding the design of our own institutions.

Three examples of topics:

Constructive value elicitation. Involving citizens in analysis requires explaining analytical choices to them, such as how to assess preferences over time or how to define “risk.” It is easy enough to write structured survey questions on these issues, and to get people to answer them. The large literature on survey anomalies, however, shows the importance of knowing exactly

which question to ask, in exactly which context. Otherwise, survey results will have little value, even when one wants people to “shoot from the hip” with their opinions. Reflecting the desire to help citizens to articulate their values on policy issues, there is growing interest in constructive elicitation, often drawing on decision analytic approaches. Suitable methods must face the behavioral reality, revealed in descriptive research showing anomalies in valuation processes (e.g., framing effects, immediate affect coloring judgments about future events, difficulty predicting adaptation). They must also incorporate a philosophy of science capable of dealing with reactive measurement. It is one thing for a private client to accept that an elicitation procedure has helped with articulating and expressing “true” values, another to accept the results of constructive elicitation as a basis for public policy. The needed research requires, at least, economics (for well-formulated policy questions) and psychology (for well-considered answers).

Decision-making competence. The division of labor between laypeople and experts often depends on their perceived competence for the job at hand. Competence is naturally context-specific. Some contexts are forgiving; others require high levels of performance. However, debates over competence often involve sweeping generalizations (e.g., the disputes between law-and-economics and behavioral-law-and-economics advocates, discussions of adolescents’ culpability for crime or reproductive rights, deliberations over paternalistic regulations). To some extent, these framings are rhetorical, seeking political advantage by citing supporting results. To some extent, they mimic the posturing of scientists, who enjoy staking out strong positions (perhaps advancing the scientific enterprise, perhaps not). However, they lack the context sensitivity needed to inform policies responsibly. They miss the opportunities for examining the external validity of research results. They fail to ask why kind of skill decision-making competence is. The research here will require both psychologists and experts in the specific contexts where performance is needed.

Behaviorally realistic risk assessment. Many complex systems (like the three focal ones for this workshop) involve the interactions of individuals, organizations, technology, and natural environments. Integrative modeling is required to anticipate and anticipate their safety (or vulnerability). The role of the social sciences in such models lags well behind that of the natural sciences and engineering. Human factors may be ignored entirely, addressed ad hoc assumptions (social variables, but not social science), or treated with computational approaches that most resemble engineering. Rarely, are the core disciplines challenged (and helped) to work up their knowledge in a form suitable for modeling (or to redirect their research agendas to this task). As a result, (possibly exaggerated) hopes are pinned on inadequate models and relevant data are not even collected. The research here requires interdisciplinary teams, including subject matter specialists who many never have thought of themselves as risk experts, and flexible, imaginative modelers.

(Two recent examples: (a) Someone called me from a (the?) consulting company doing a risk analysis for plague in the hands of terrorists; he wanted me to come to MacLean for a day to provide numbers for the human factors. (b) The Executive Director of the National Association of City and County Health Organizations told me that no one knew how many calls public health organizations had fielded during the anthrax crisis, even though those calls came close to paralyzing the systems and undermining their credibility. These capacity issues are, presumably, among others outside any systematic analysis.)

Geoffrey Heal

Notes on the economic issues raised by national security policies.

From an economic perspective there are three issues that seem central to security issues:

1. Uncertainty and risk
2. Public goods and
3. External effects.

Uncertainty and risks

The centrality of risk and uncertainty is self-evident. A key feature of the uncertainty here is that it is hard to describe. Speaking in terms of Savage's axioms, it is clear that we cannot describe the entire state space. There are things that might happen that we cannot anticipate – who, for example, a year ago would have given any thought to the possibility of two large jets being hijacked and crashed into the Twin Towers simultaneously? There are obviously other unanticipated events waiting to happen. Formally, we have a state space that is only partially described. And for the part of the space that we can describe, which presumably consists of events like bombings and bio-weapon attacks, the probabilities are certainly not well-known. We almost certainly do not have anything that could be described as a probability distribution over them. All of this makes the uncertainty hard to describe and hard to model. There are similarities with the area of climate change – there we do not understand fully the possible consequences of climate change and we cannot assign probabilities to the events that we can describe. (Heal and Kristrom

<http://www.gsb.columbia.edu/faculty/gheal/General%20Interest%20Papers/climatereview.pdf>)

The description of rational behavior in these circumstances is challenging, as is the evaluation of policy options. We need research on these issues.

The complexity of this area clearly raises questions about how individuals will react to it, which must be central to any policy analysis, but there are others on the team who are better placed than I to speak to this.

The complexity of the uncertainties that we face also mandates an interdisciplinary approach here: as an economist I cannot evaluate the possible types and outcomes of terrorist attacks. This must involve experts in engineering, weapons technology and anti-terrorist measures.

What is it worth paying for better security? This is a natural economic question. The answer will depend on the possible losses and on the population's attitude towards risk. In the case of poorly-described risks such as we are addressing here, we have little analysis of how to evaluate what people are willing to pay for risk avoidance. This is an interesting research topic.

Public goods

Security is a public good: if the community is more secure for one then it is more secure for all, at least to some degree. However, it is an unusual public good: it has to be provided not only by the public sector but also to some degree by all citizens and corporations. As we have seen, private corporations such as airlines can contribute to security or to its absence.

Telecommunications companies affect our security by the degree of redundancy that they choose in their equipment. Individual citizens affect it by their vigilance. So unlike the traditional public goods that are provided by the state, this public good is necessarily provided by a partnership between the public and private sectors. What are the terms of this partnership?

There is an analogy here with the environmental area again. Climate change is a public bad that is provided not by government action but by the actions of all individuals and corporations who burn fossil fuels. In both cases we have what have been called privately produced public goods. A standard method of regulating the provision of these in the environmental case is through cap and trade systems, such as the 1990 amendments to the Clean Air Act (see <http://www.gsb.columbia.edu/faculty/gheal/EnvironmentalEconomicsPapers/pw-94-03.pdf>).

What is the analog to cap and trade in the security area? Is there a role for a tradable risk permit? In general we do not expect market mechanisms to manage the provision of public goods efficiently, but there are exceptions. One is the use of cap and trade systems with privately produced public goods. Another is the bundling of public and private goods in cases in which the public good enhances the willingness of buyers to pay for the private good: in cases such as this the seller may have an incentive to provide the public good at an efficient level. Although airlines have traditionally bundled security with the private good or transportation, it seems in retrospect that they have not provided this public good at an efficient level. This is predictable as a function of the industrial structure of the airline industry (see http://www.gsb.columbia.edu/faculty/gheal/EnvironmentalEconomicsPapers/pubgoodspropvalue_s.pdf). What are the implications of these results for the use of market-based systems to provide some aspects of security?

External effects

An important aspect of security problems is that there is an element of interdependence between the security levels that agents can attain. My security may depend on the measures that you take, and vice versa. This introduces externalities to the problem. A sharp example comes from airline security: poor security procedures at one airport or by one airline may allow passengers with criminal intent or explosive baggage into the air transportation system, so that they can transfer or be transferred to other airlines and airports. This interdependence of risks can reduce the incentives that any agent has to manage these risks (see <http://www.gsb.columbia.edu/faculty/gheal/General%20Interest%20Papers/onlydieonce.pdf>). How do we design institutions that react to these externalities and the interdependence that they introduce?

Stephen Hilgartner

The Social Dimensions of Expert Knowledge about Risk

I. General Problem Area

Studies of the social dimensions of scientific and technical knowledge about risk should be an important component of the NSF's program on risk analysis and decision making. In many risk areas, decision making involves complex negotiations about the strength of technical evidence—often in a context of scientific uncertainty, disagreement about the credibility of different types of expertise, and struggle among a variety of stakeholders. Indeed, the social and political challenges of creating credible sources of expertise and credible knowledge often rank among the significant obstacles to making sound policy about risk. For this reason, society's ability to address risks depends not only on sophisticated analytic methods and detailed information, but also on understanding the production and reception of expert knowledge; the politics of integrating knowledge and producing credible expert advice; and the problems of legitimacy that social institutions face in risk management.

To address these problems, careful empirical research—using a variety of ethnographic, sociological, historical, and comparative approaches—is needed in several areas:

(1) The social processes that shape how expert knowledge about risk is produced, contested, evaluated, and utilized.

Experts play a central role in creating knowledge in risk analysis, and they bring particular theories, training, techniques, and practices to their work. These outlooks shape how problems are framed, what data are collected, and what management strategies are considered. In different technical communities, different cultures of risk management are found. For example, hospitals bring an individual orientation to the management of medical errors that differs greatly from the approach to error taken in, for example, commercial aviation. Such differences can profoundly influence how risks are defined and what is done about them. Beyond framing problems and generating data, experts are involved in evaluating evidence that, in many cases, everyone agrees is incomplete, uncertain, or contested. Reaching collective judgments about the credibility of knowledge claims has an important social dimension, since diverse communities may develop different views about when one can safely extrapolate from test conditions to real world conditions, treat computer simulations as adequate models of actual phenomena, or make future predictions based on historical experience. Negotiations and struggles over the credibility of knowledge are a salient feature of most risk controversies. Detailed studies of the social processes through which expert knowledge is produced, contested, and certified can illuminate this dimension of problems of risk.

(2) The framing of risks in complex sociotechnical systems.

Many significant risks—from global climate change, to transportation accidents, to computer security—are embedded in large sociotechnical systems, networks that weave together a variety of machines, people, procedures, laws, and other components. These systems typically span

organizational boundaries, and, in many cases, were shaped not by a centralized design plan but through an evolutionary process, unfolding over decades and involving many spatially and socially distributed actors. Entities that contribute to risk—or “risk objects”—are spread throughout these systems. For example, the risk objects that contribute to motor vehicle accidents include drivers, cars, roads, weather conditions, alcoholic beverages, bars accessible only by automobile, and so forth. Risk management is by no means a post-hoc “add-on” to technology; controls aimed at reducing risk are woven into the basic fabric of large technological systems. For example, at every step in the long journey from the oil field to the spark plug, fuel flows through a network of pipelines and procedures aimed at preventing premature combustion.

In complex systems, it is often difficult to unambiguously identify important sources of risk: many risk objects are present, causation is typically multifactorial, and proximate causes may be neither the most fundamental causes or nor the most susceptible to control. Such ambiguities provide a strong motivation for conducting studies of how experts frame problems of risks in complex sociotechnical systems. How do they select particular objects for research attention and at the expense of what alternatives? How are neglected objects identified, studied, and made salient to policymakers? In large, distributed systems, do jurisdictional boundaries (of government agencies, scientific disciplines, and other knowledge producers) influence the kinds of data collected and controls contemplated? Studies that examine how knowledge about sociotechnical systems is produced may be able to improve risk management strategies.

(3) The challenge of producing credible expert advice on risk.

The challenges of producing credible science advice and institutions for creating it looms large in risk policy. In controversial areas, charges of bias, incompetence, and conflict of interest appear with some regularity, and complaints that science advisors have made inappropriate recommendations in light of the strength (or weakness) of the data are nearly ubiquitous. Critics of science advice and the institutions that produce it are often skilled at exposing simplifying assumptions, gaps in evidence, and other sources of contingency in order to undermine reports. The process of producing advice also often comes under attack; for example, for failing to ensure appropriate representation of all relevant domains of expertise, stakeholder groups, or, in international settings, nations or regions. The institutions that produce science advice, for their part, work to protect themselves against such attacks, and the credibility of advice and advisors often emerges from a dynamic interaction. Studies of science advice and institutions involved in producing authoritative risk assessments and management plans are needed. Because advisory mechanisms differ greatly among institutions (both within the United States and cross-nationally) comparative research on advisory systems are likely to be especially informative. In addition, there is a need for additional studies of how advisory institutions in different polities have approached problems of consensus-building and structuring public participation in the deliberative process.

(4) The problem of risk and the legitimacy of societal institutions.

Studies are also needed of the role of risk management in establishing trust in societal institutions. In many domains, risk analysis has provided a means to help depoliticize controversies about new technology, converting volatile mixtures of technical uncertainties, science fiction scenarios, critiques of corporate power, and misgivings about hubris and slippery-

slopes into categories amenable to bureaucratic management. In other cases, such as the BSE debate in Britain, failures of risk management and risk communication have significantly undermined the legitimacy of public institutions. Research is needed that takes a critical look at the politics of risk management, assessing its strategic use and considering the risks to political legitimacy associated with its failure.

II. Specific Research Challenges: Illustrative Examples

The social organization of data production. How does the social organization of data production about risk affect the kinds of information available for risk assessment? Do the organizational routines that generate official statistics shape the possibilities for risk assessment? Do systematic gaps in knowledge result from patterns of data collection in government and corporate bureaucracies? Do jurisdictional boundaries constrain the scope of risk analysis in some domains? Do restrictions on access to data (e.g., for proprietary reasons) constrain risk analysis in important ways?

Historical studies of risk assessment methods. How do techniques for estimating risk become established and win credibility? How do communities of practice form around particular risk assessment methods, and how do these communities build shared assumptions, tools, and standards? Can one find regularities in the historical process of developing risk assessment methodologies that shed light on the challenges of building credible assessment methods?

Neglected problems and new strategies. Can studies of diverse cultures of risk management help identify neglected sources of risk or approaches to risk management?

Studies of attribution of responsibility for disasters. In the aftermath of accidents and disasters, such as Bhopal, the Challenger explosion, or the BSE crisis, how do investigators and other interested parties struggle over causal responsibility, moral responsibility, and the responsibility to address problems? Are these various forms of responsibility attributed to people, machines, procedures, organizations, laws, nature, etc.? When—to take the extremes—is blame localized to a single point in the system, and when is it attributed to the system as a whole? When are failures treated as normal features of the operation of a system and when are they considered abnormal intrusions into it? Through what processes do such attributions (at least some of the time) become conventional wisdom and influence future action?

Studies of negotiation about credibility. How do experts from different technical fields assess, debate, and negotiate about the relative credibility of knowledge about risk produced using diverse methods? How can institutions develop mechanisms for improving the integration of information from diverse sources of expertise?

Changes in risk management paradigms. How do paradigm shifts in risk management (such as changes in emphasis from controlling individual errors to engineering safety into systems) occur in particular technical domains? Historical studies of past shifts and studies of contemporary risk management paradigms (especially those that are being challenged) are needed.

Comparative studies of institutions. How have different national governments and transnational institutions approached the challenge of producing sound advice and legitimate decisions about similar risks?

III. Promoting Future Research and Cross-Disciplinary Collaboration

I would suggest that to stimulate research on such topics, studies of the social dimensions of expert knowledge about risk should be explicitly designated as among the priorities for research on risk analysis.

To expand the number of scholars with the requisite training to address these issues, expansion of graduate education is needed, perhaps using the NSF's IGERT mechanism.

To encourage collaboration between natural scientists and social scientists in this domain, several directions merit consideration:

- A multidimensional workshop dedicated to these issues would be worthwhile.
- One promising mechanism would be postdoctoral fellowships to natural scientists who seek to develop expertise in science & technology studies. Fellowships for social scientists who wish to develop expertise in risk issues would also be useful.

Hank Jenkins-Smith

I would offer four areas of focus in risk research that, in my view, merit serious attention. Useful progress in any of these problems would require interdisciplinary research.

(1) As noted in the workshop materials, a central objective is to identify obstacles to effective use of existing research findings and tools, and to suggest ways of overcoming these obstacles. In my view, one such obstacle is the close and persistent entanglement of risk research with policy advocacy.

I am concerned that the development of risk analysis and risk-related research has become so infused with policy advocacy and the promotion of underlying value dispositions that the results of risk research will become seen by policy makers as mere window-dressing for other ends. This entanglement can be seen in a range of developments, including the bounding and emphasis of risk assessments by federal and state agencies (e.g., DOE's unwillingness to include risk perception-based impacts in its analysis of the effects of the Yucca Mountain high-level nuclear waste repository, or the State of Nevada's apparent mobilization of risk research to amplify the case against the repository) and efforts to use potential risk to publics as a means to champion alterations in collective decision-making structures (e.g., the implementation of site specific advisory boards, "dialogue groups", and other mechanisms that single out organized interest groups for representation).

As risk scholars, we owe it to decision makers to be clear about the nature of the systematic connections between risk research findings and underlying values, policy preferences, and value structures. At the same time, policy makers committed to the public interest may (sometimes correctly) be reluctant to support risk research if it appears to be a carrier for considerations beyond the scope of the event or policy in question.

As risk researchers, we need to learn from the ways in which risk research has been employed in major policy debates. I would encourage the NSF to sponsor research on the ways in which risk research has been employed in major policy domains (nuclear waste, ballistic missile defense, acid rain, global climate change, and others). The research should focus on the different uses to which risk research has been put, and the ways in which those uses have affected the utility and credibility of the research findings. The intent would be to understand the "obstacles to effective use" that have emerged and to develop means to overcome these obstacles.

(2) The NSF, other government agencies, and NGO's spend substantial resources to "educate the public" regarding environmental and consumer risks. Despite these expenditures, relatively little is known about if, how and when individuals use this kind of information to update their understandings of the risks with which they are confronted.

Research should focus on the ways in which people can acquire and process information about risks when it is made available to them. What are the attributes of those individuals who are willing to "invest" in acquiring new information? How does that investment affect knowledge

(i.e., when can new information alter prior beliefs)? In part this research would focus on risk communication, but I would urge that the focus should be more on the basic understanding of the processes by which members of the public seek out, utilize and process information about potential hazards.

This research faces several challenges, the most serious of which is to measure and model investment in information acquisition and the changes in beliefs and preferences that such investment might bring about. These processes are almost certainly endogenous, in that some of the attributes that lead a person to seek new information (e.g., a belief that a phenomenon is hazardous) will also be influenced by the acquisition of the information. A second challenge concerns the external validity of the research, in that the information needs to be provided in a way that reasonably approximates the manner in which such transmissions actually occur, providing a reasonably balanced and thorough treatment of the potential hazard, and encompassing the relevant uncertainty on the matter.

As a further step, I would want to see investigations of how knowledge and subjective certainty influence willingness to act when individuals are faced with costly actions to address the potential hazard. When is the effect directional (e.g., when does greater knowledge increase the perceived risk and willingness to take costly actions), and when does it affect the predictability of choice (introducing a heteroscedastic element to modeling efforts)?

The objective of this line of research would be to understand how it is that information about hazards is acquired, accepted (or rejected), and factored into choices about potential hazards. The practical implications for how government agencies provide information about risks – such as terrorism, global climate change, nuclear waste transport, or second-hand tobacco smoke – are substantial.

(3) I would urge NSF to fund research that builds understanding of the ways in which risk experts and scientists manage uncertainty as they make the shift from identifying the most “scientifically supportable” assessment of a hazard to the assessment that should guide risk standard setting. Several examples:

- a. What dose-response relationship between radiation exposure and cancer incidence is seen as best supported by the current body of scientific findings; and what dose-response relationship should be assumed when specifying health standards for radiation exposure?
- b. What relationship between greenhouse gas emissions and global temperature patterns is understood to be best supported by the body of scientific evidence; and what relationship should be assumed when setting emission reduction objectives?

For either example, how does the level of uncertainty influence the shift from the “scientifically justified” relationship to one that is preferred for regulatory standard setting? The step from identifying relationships that are most consistent with scientific findings to those that should underlie standards is critical, and infuses such debates as those concerning the Delaney Clause or the “precautionary principle”. To the extent that the basis for taking this step among the scientific and technical communities that advise policy makers remains opaque, a crucial component of risk policy making cannot be examined and evaluated.

(4) The relationship between stated risk perceptions and intended behaviors as measured in surveys of the public, and the actual behaviors and choices of individuals, needs to be better understood. One useful avenue would be to combine survey research on perceived risks and contingent valuation for risk reduction with measures of actual behavior, such as hedonic analyses of real estate decisions where a risk component is an attribute of the property or experimental research wherein choices are more than hypothetical. While some work of this kind has been attempted (e.g., Gawande and Jenkins-Smith, 2001; McClelland, Schultze, and Hurd, 1990), the limited nature of the data connecting risk perceptions or stated intentions with real transactions has precluded firm findings. One of the main challenges facing this line of research would be to develop overlapping datasets that would permit robust assessments of the relationship between survey-based findings and those derived from choices and transactions.

Ralph L. Keeney

Thoughts on an NSF Risk Program

Objectives of the Research

The appropriate problem areas for research naturally depend on the objectives to be achieved by that research. In the workshop description, guidance about objectives is given:

- Apply risk analysis and decision-making tools more effectively to address societal problems
- Identify obstacles to using these tools and related knowledge
- Suggest ways to overcome the obstacles.

It is important to recognize that the objectives above are all means objectives. I think that the fundamental objectives that are the real basis for our interest in risk problems and in this risk program are the following:

- Minimize loss of life and detrimental health effects
- Minimize disruption of society and its individuals (e.g. fear, despair, loss of jobs)
- Minimize detrimental impacts to the natural environment
- Minimize economic costs, both direct and indirect.

In general, the intent is to maximize the quality of life and minimize any detrimental impacts on the quality of life of citizens of our country.

Focus of the Research

I think that the adequacy of our methodologies is not the relative weak point for achieving these objectives, but that our major weaknesses have to do with effectively applying what we know and effectively communicating the knowledge we have and insights that we can get from applying our knowledge.

From my perspective, the implication is that much of the research should be applied (e.g. how to conduct analyses, how to use analyses, policy analyses) rather than theoretical or methodological.

We Should Strive for a Very Big Recognized Success

As part of the research program, I think we should have one project or a focus of several projects that should provide a very big recognized success. For example, if we could make the statement that the \$50 million invested in the program saved a minimum of 10,000 lives, this should make an impact on people in Washington and the citizens of our country. This would greatly facilitate the support and receptivity towards other more subtle and more long-term contributions that we could make but still would be very important.

A candidate for a big success would be to get a policy set for the value of a statistical life or the value of a year of life that is used for many government investments. We know there are governmental programs that spend hundreds of millions of dollars and more to save one statistical life. We also know there are many potential programs where statistical lives can be saved for a few million dollars at most. If a policy led to understanding the implications of misallocation of funds across programs and further lead to the better allocation of such funds, we could easily save 10,000 lives *per year*.

If we did a survey of people in Congress or of lay public and asked them whether they would prefer five beautiful new risk methodologies or saving 10,000 American lives per year for the same amount of investment, I think the preference is obvious. Let's have the commitment to make that choice as part of the research program.

Include Values as Part of the Research

There are some who may think that values aren't science and should not be studied as part of a scientific research program. I would ask, why is it that we care about risk at all or this proposed risk program? The answer is our values. We want better consequences rather than worse consequences, where better and worse are characterized by how well we meet the fundamental objectives indicated above. Without values, there is no reason to spend any time on any risk problem. Also, without thinking about the values we hope to achieve by dealing with risk problems, one couldn't possibly address them in any reasonable or responsible manner.

The amount of research funds spent on developing methodologies to relate facts and on gathering factual information and making sense of it is orders of magnitude more than the research funds spent on understanding values that drive our interest in risk problems. A reasonable investment in the value aspects of risk problems should therefore have a potentially high payoff.

Possible Topics of Research

The following are a list of some potential topics of research with no prioritization implied.

- Educate legislators and policy makers that tradeoffs must be made and always have been made and that explicit thinking about such tradeoffs can be useful.
- Develop and implement some reasonable ways to obtain tradeoffs that would be defensible and could be consistently used.

- Get a policy set for the value of a statistical life or the value of a year of life that should be used for many government investments. We could "save" a lot of lives if this happened.
- Develop general guidelines for how to conduct risk and decision analyses. Examples include evaluating all reasonable alternatives, doing sensitivity analyses, and addressing explicitly all the features that make the problem complex.
- Develop possible general guidelines for how to present and interpret risk analysis (and models in general). Examples include recognizing uncertainties, presenting sensitivity analyses, and stating that no analysis yields the definitive "answer" to a policy problem.
- Understand all the reasons why risk analysis and decision-making methods are not liked and/or not used by the public, legislators, the media, and the legal system. Develop ways to lessen such problems.
- Develop procedures to help create a rich array of alternatives to address specific risk management decisions.
- Investigate ways that it would become politically correct for politicians to clearly address their viewpoints about the complex issues inherent in risky decisions of concern. They should for instance, recognize significant uncertainties and value tradeoffs and perhaps even express their judgment on these.
- For the research program, develop a strategy and a way to follow it that enhances the chances that the program objectives are achieved. If it is different people handing out small research projects to numerous people who apply for funds, our chances are less than they need to be.
- Develop methodologies to incorporate civil liberties and rights as objectives in decisions that may be faced regarding terrorism. We want ways to measure the degrees to which civil liberties and rights are limited or expanded by various acts. We need to have ways to indicate which groups of people or individuals have their liberties and rights affected in different ways and address the equity concerns of those. We need to understand the tradeoffs of civil liberties and rights versus the right to safety, life, and the pursuit of happiness.

Paul R. Kleindorfer

Risk Management Challenges

Potential Problem Areas of General Import

The discussion provided as part of the Workshop background materials seems appropriate as a framing of the key problem areas of likely interest to NSF. This discussion mentions:

- The traditional disciplines/foundations of risk management (assessment, perception, etc.) as well as decision sciences foundations (choice, ambiguity, legitimation, etc.)
- Problem areas of homeland security, natural and technological hazards, environmental and international health issues.

In terms of problem areas of general import for risk analysis and decision sciences, let me mention two broad areas for consideration and discussion at the Workshop: (1) Environmental, Health and Safety Regulation and Practices and (2) Risk Management in the e-Economy. I will note under each of these areas what I think the problem is and what some of the applied and basic research challenges are.

Environmental, Health and Safety (EH&S) Regulation and Practices

The focus of this problem is on public policy, regulation and management practices related to EH&S. This includes areas such as climate change and terrorism, as well as traditional areas for risk analysis such as nuclear safety and plant safety.

One arena in which I have been working, with many others, in the past few years has been risk management in industrial facilities, with a primary focus in my case on the chemical industry. There the basic research context has been characterized by the availability of new data sources and new codified management systems. An example of the former is the accident history database under section 112(r) of the Clean Air Act Amendments. An example of the latter is ISO 14000 or EMAS standards that specify a set of auditable practices that are supposed to embody “good” practices in relation to managing a firm’s EH&S performance. The Wharton Risk Center project on Accident Epidemiology is an example of what can be accomplished here and what may augur the type of research NSF may wish to stimulate going forward. Using the tools of medical epidemiology, data on 15,000 chemical facilities were analyzed, together with related information on the financial characteristics of parent companies of these facilities, characteristics of the surrounding community as captured in census data, and a number of details on the facility and the nature of the chemicals used there. The general idea was to determine from historical data which plants appeared to be “sick” (likely to have a severe accident) and which ones were “healthy”. Coupling this data analysis with an assessment of regulations applicable to such facilities and to management practices embodied in certification procedures

such as ISO 14000 leads to the possibility of determining robust precursors of good practice, in the specific arena of EH&S performance for industrial facilities as well as generally for risk regulation. The implications of this sort of analysis for public policy, regulatory monitoring and enforcement, and managerial practice are potentially significant.

Research challenges here include a range of economic and regulatory questions, including methodological issues related to evaluation, choice under conditions of complexity and ambiguity, use and legitimation of model-based assessments, new data analysis methods such as data mining and AI techniques, the performance of performance-based regulation and other familiar topics from the literature on the economics of risk regulation. What is new is the scope and detail of available data to study these issues.

A related area to that of EH&S performance of industrial facilities is the extent to which risk management practices utilized for EH&S effects may also be extended to address the problems of terrorism. The Wharton Risk Center has since then done a number of studies on this matter since 9/11. The ensuing dialog has been complicated, as it involves a reexamination of the basic tenets of economic behavior:

1. Property rights: If an attack did occur, who owns the liabilities for damages? How large do the losses have to become before what might be considered a private loss becomes a social responsibility?
2. Mitigation and Prudent Precautionary Behavior: How should standards be developed (by the facility owner, the trade association, the government, ...) and who should pay for improving site security?
3. Models of Man: Should terrorism and associated economic responses be thought about primarily in terms motivated by self-interest, or should a different calculus be used, perhaps based on industry self-regulation, or on social norms to respond to this threat?
4. Risk Management: How should the traditional model of risk management, based on randomly triggered events, be modified to accommodate triggers arising from purposeful and malevolent agents? What prioritization procedures are implied by the required modifications?
5. How should insurance markets and risk management professionals respond to these threats in terms of coverage, premiums, pricing of professional services, and so forth? What is the role of private insurance in this arena, given the large role government plays in detection and deterrence?
6. How should public-private partnerships be crafted in the face of very ambiguous vulnerabilities and risks?

In my view, we have not even begun to sort out the above issues. The basic problem is that the debate until now has been colored by political ideologies and not by objective research on the consequences of one or another strategy. This results in part because we are still confused about what to make of “terrorism” as a phenomenon, at least in the United States. It also is the result

of a confusion about what private markets can and cannot do, and what the appropriate role of national and international institutions is. Similar problems are apparent in the area of global climate change and the U.S. non-response to the call to endorse the Kyoto Protocols for reducing greenhouse gases.

The above brief thoughts convey, I trust, my conviction that basic research has an immensely important and continuing role here to inquire as to what courses of action are open to us in assessing and managing EH&S risks, including those arising from terrorism, and what their likely fit and performance would be when coupled with existing economic, political and social structures.

Risk Management and the e-Economy

New institutional forms are arising in the e-economy that are using new methods for integration of business activities, from investment, to marketing to supply chain management and beyond product use to environmental impacts of ultimate recovery or disposal of products. These institutional forms are founded on time-honored concepts of self-interest and property rights, but the volume and speed of data and informational interactions in the new economy pose significant new opportunities and challenges for research in risk management and the decision sciences. I will illustrate this briefly in one of the areas I have been exploring with others in the Wharton Risk Center, the use of new hedge and derivative instruments, coupled with auctions and other economic mechanisms to coordinate economic activity via B2B and B2C exchanges. To keep my problem to within manageable proportions, I focus only on B2B exchanges for capital intensive industries, but similar research programs for other aspects of risk management in the e-Economy should be evident.

Arguably, the central problem addressed by e-Business is better coordination of supply and demand, including price discovery and reduction of transactions costs of buyer-seller interactions. In capital-intensive industries like chemicals and steel, the out-of-pocket costs of excess capacity and the opportunity costs of underutilized capacity have been important factors driving the growth of exchanges for improving demand and supply coordination through e-Business platforms. The emerging reality of such exchanges is, however, quite complex in terms of product range, protocols for joining the exchange, rules for operations and settlement and so forth. And we are only beginning to understand some of the challenges of designing such institutions, as the Enron debacle and recent broadband disasters show.

A particularly interesting development in this arena has been the parallel development in B2B markets of long-term and short-term markets for capacity and output, accompanied by a range of exotic options and forwards as the basic mechanisms supporting transactions. A central feature of B2B is that, especially for capital-intensive industries like chemicals, steel and electric power, contracting needs to take place well in advance of actual delivery. Failure to do so for a non-scalable technology is a recipe for last-minute confusion and huge excess costs. This has given rise to a general recognition that most of a plant's output in such industries should be contracted for well in advance. However, there is still a very important role for short-term fine tuning of capacity and output to contract for, say, the last 10\% of a plant's output or a customer's requirements. Doing so requires a conceptual framework, congenial to e-Business, that allows

contracting to take place at various points of time, constrained by various commitment and delivery options and flexibilities, and mediated by electronic markets where these are feasible. The research challenges here are significant, in both scope and potential impact.

To illustrate, consider the electric power industry. The deregulation of the sector has led to an explosion of activity by financial and energy intermediaries, offering both futures and options products. These exist through the New York Mercantile Exchange in the U.S. and through other exchanges throughout Europe, Latin America and Asia. The information being generated from these long-term and short-term contract exchanges is further giving rise to a new range of risk hedging instruments, such as the weather derivatives recently introduced in the electric power industry. These products provide various rights and obligations to buyers and sellers for the purchase of energy delivered at specific places and at specific times. These products are clearly substitutes for equivalent deliveries in the associated spot markets. The use of these instruments in the natural gas markets and in electric power has literally transformed the entire industry in the last decade. And behind the trading divisions of nearly all major energy companies is a growing recognition of the premier place that research plays in enabling these companies to use these new risk management tools appropriately and profitably. The same general trends are visible in many other industries, including semiconductors, commodity chemicals, plastics and other types of capital-intensive goods. The resulting market structure has stimulated an explosive growth in financial and contractual innovations, and almost all of these await both the market and research to sort out which instruments, exchanges and contract forms can be expected to survive and provide the value added awaited from these new approaches to risk management.

Conclusion

In both of the above arenas, the challenge will be to determine what mix of problem specific attributes, cross-disciplinary and basic disciplinary issues (from risk management and decision sciences) to emphasize in NSF programs going forward. I look forward to the discussion at the Workshop to help to clarify this.

Howard Kunreuther

1. What problem areas would be appropriate for consideration in a new integrated research program on risk analysis and decision science?

- Linking science with policy for dealing in extreme events. There is a real opportunity for research to be undertaken between the engineering and scientific communities and the social sciences to deal with risk-related issues. The area of extreme events raises a set of special challenges because, by definition, there are limited historical data and one therefore has to rely on scientific risk assessments to supplement past experience.
- Challenges associated with mitigation and loss protection for dealing with a wide variety of risks including terrorism. Protective measures normally involve a tradeoff between incurring a certain upfront cost associated with an investment in return for uncertain benefits over a potentially long time-horizon in the form of reduced losses should a disaster or accident occur. There are a number of behavioral issues associated with how individuals, groups, organizations and society make these tradeoff and challenges for improving their decision-making processes.
- Appropriate public-private partnerships for dealing with catastrophic risk. Following some severe natural disasters in the 1990s and the September 11th World Trade Center terrorist attack it is now becoming increasingly clear that one needs research on the appropriate types of private-public partnerships that are likely to work in practice. One cannot expect private market forces to provide adequate economic incentives to encourage individuals and firms to invest in protection nor can one rely alone on governmental agencies to enforce regulations and standards. Both sectors need to work together to deal with the problems of managing catastrophic risks.

2. For the above problem areas, what are the most significant research challenges that we should consider?

- Research on understanding the bases for expert judgments, why they differ from laypersons for a variety of risks and what ways these two groups can better appreciate each other in order to develop more meaningful risk management strategies. For example, what are better ways of characterizing experts' views on uncertainty and providing data on this uncertainty to the lay public who have to make decisions?
- Research on understanding the role that ambiguity and uncertainty play in individuals' and organizational decision-making processes. For example, why have many insurers discontinued terrorism coverage following Sept. 11th? Why has the government spent so much money on *reassurance* (as well as reinsurance) following Sept. 11th?

- Research characterizing the nature of individuals' group and organizational decision process with respect to collecting and processing information associated with extreme events. We need to gain more insight as to when individuals and groups utilize analytic approaches in making choices and their interaction with affect-related factors (e.g. worry, stress, fear). Both play a role in the decision-making process. Research also needs to be undertaken on what criteria should be used to judge how good a decision is likely to be?
- Based on our understanding of decision making for different type of risks, research is needed to improve the risk management process. There is a need for research on the expected benefits and costs of different types of public-private partnerships for specific risk-related problems. Analyses can include presentation of information, types of economic incentives (e.g. subsidies, fines), collective choice approaches and social norms as well as the use of regulations and standards.
- Research could be undertaken on the types of public-private partnerships for dealing with the following types of risk-related problems:

Dealing with misperceptions and misestimates of certain risks due to systematic biases (e.g. availability) or simplified decision rules (e.g. "it "won't happen to me")

Dealing with negative externalities associated with catastrophic accidents (e.g. decrease in property values; contamination of one individual, group, and/or organization to others due to interdependent security problems)

Inability for the public sector to enforce certain regulations (e.g. limited personnel and expertise to undertake adequate third party inspections)

- Research is needed on understanding the role that decision analysis and benefit-cost analysis can play in evaluating alternative strategies for dealing with different types of risk-related events. Some of the specific issues that could be examined are:
 - How to deal with differences in values, agendas and perceptions of risks of different stakeholders
 - Tradeoffs between distributional (equity) and efficiency (optimal resource allocation) criteria for choosing between alternatives and/or projects
 - Dealing with short and long-term issues such as taking into account future generations and what social discount rate to use
 - Role that uncertainty and ambiguity play in determining choices between alternatives
- Research on what it takes to establish trust when one constructs risk management strategies of extreme events both before the event happens (e.g. to get people to adopt loss reduction measures) and after an accident or disaster (e.g. rebuilding and recovery strategies)

- Research on phenomena that are not well understood in characterizing people's behavior toward risks. Three that come to mind are:

Tipping behavior---actions of one individual or a small group may cause others to change their behavior;

Role of social norms and coordinating mechanisms in improving risk-related decisions (e.g. organizational associations and their roles);

Impact of short-time horizons and/or high discount rates in inhibiting certain behavior and how one can overcome them

- What are some of the challenges in implementing risk management strategies for extreme events and how can address these questions?

3. How can an NSF initiative most effectively promote future research in risk analysis and decision sciences that brings together the natural and social sciences?

- There is an important role that Risk Centers can play in bringing different disciplines together and undertaking meaningful research linking science with policy. I feel there is a need for problem foci in such Centers in order to help reduce the communication gap between different disciplines
- There is an opportunity for one to engage in creative partnerships between governmental agencies, industrial firms and the research communities in developing strategies that have a chance of being implemented. NSF may want to partner with specific federal agencies (e.g. EPA, OSH, Office of Homeland Security) in supporting research that promises to provide social benefit.
- Following Sept. 11th there is an opportunity for this NSF initiative to encourage research in new areas (e.g. developing strategies for dealing with terrorism) by utilizing techniques and approaches that have been successful in other areas (e.g. natural disasters and technologic risks)
- There is a need for developing data sources that can be shared among principal investigators (PIs) supported by NSF. In fact, some type of joint workshop may be desirable where NSF bring together those supported by this initiative to facilitate interchange of data and ideas.

What is the appropriate balance between having this initiative devoted entirely to supporting Risk Centers and entirely supporting individual PIs? What types of synergy should be encouraged between different research projects and what funding and logistical mechanisms could be set up by NSF to encourage this?

Peter R. Orszag

What problem areas would be appropriate for consideration in a new integrated research program on risk analysis and decision-making?

An integrated research program on risk analysis and decision making would be particularly beneficial to policy-makers struggling with homeland security issues. As one senior policy-maker in the homeland security arena recently put it, an insufficient stock of “off-the-shelf” research on homeland security exists to inform policy-making. The relatively modest amount of relevant research raises the possibility of “low-hanging fruit” – research that may be able to draw upon lessons from other areas to improve policy-making significantly at relatively low cost. Furthermore, as the introductory note for the conference emphasizes, homeland security research would likely require a multi-disciplinary approach: “The variety of potential targets, the political nature of the underlying factors causing terrorist activity, and the economic consequences of security measures mean that homeland security necessarily involves multiple disciplines in both the natural and social sciences.”

For the above problem areas, what are the most significant research challenges that we should consider?

As several co-authors and I argued in a recent Brookings volume,¹ some of the most challenging issues in homeland security involve government intervention in private-market settings. When should the government intervene to require, finance, or provide security at private-sector sites? How should it intervene?

Industry groups will invariably claim that voluntary efforts would provide adequate levels of security in private-sector settings. The Brookings volume argues, however, that there are at least six potential justifications for government intervention:

- First, a significant terrorist attack undermines the nation’s sovereignty, just as an invasion of the nation’s territory by enemy armed forces would. Firms will not generally take this cost into account in evaluating security measures.
- Second, a more specific negative externality exists with regard to *inputs* into terrorism. For example, loose security at a chemical facility can provide terrorists with the materials they need for an attack on other targets.
- A third potential motivation for government intervention involves information costs. For example, each individual could individually conduct biological safety tests on the food that he or she was about to consume. The information costs associated with that type of system, however, make it much less attractive than a system of food safety regulation.

¹ Michael O’Hanlon, Peter Orszag, Ivo Daalder, Mac Destler, David Gunter, Robert Litan, and James Steinberg, *Protecting the American Homeland: A Preliminary Analysis* (Brookings Institution Press: 2002).

- The fourth justification for government intervention is that corporate and individual financial exposures to the losses from a major terrorist attack are inherently limited by the bankruptcy laws. The incentive to invest in protective measures is therefore muted.
- The fifth justification for government intervention is that the private sector may expect the government to bail it out should a terrorist attack occur. Such expectations create a moral hazard problem.
- The final justification for government intervention involves incomplete markets. The most relevant examples involve imperfections in capital and insurance markets. In addition, certain types of activities may require large-scale coordination, which may be possible but difficult to achieve without governmental intervention.

The importance of these six factors varies from situation to situation, and the degree (as well as existence) of government intervention should therefore also vary by circumstance. For example, consider the difference between security at a shopping mall and security at a chemical facility. Poor security at a mall typically does not pose a risk to as many lives as poor security at a chemical facility. Thus security regulations for chemical plants may make sense, even if they don't for shopping malls. The point is that some threshold must be set for intervention.

The standard we set for focusing homeland security efforts is the prevention of those attacks that would involve the loss of thousands of lives, the interruption of crucial social and economic services for millions of people for an extended period, or significant damage to the reputation of the United States in the global community. To be sure, protecting against such extremely costly acts of terrorism may displace attacks toward less damaging scenarios, involving fewer people or less important infrastructure. Since it is impossible to eliminate all potential terrorist activity, however, such displacement of terrorism toward significantly less damaging attacks is a necessary consequence of providing better prevention and protection against the most damaging attacks. In other words, some displacement is unavoidable, and it is notably better to displace terrorist activity toward significantly less damaging attacks than to displace it from one catastrophic area to another.

Crucial research questions that follow from the Brookings preliminary analysis include, among others:

- Most broadly, the relative importance of the factors delineated above in undermining the efficiency of the private sector in protecting itself in the absence of government intervention;
- The interplay between positive and negative externalities in private-sector security activities (risk-shifting from one potential target to another imposes a negative externality from private security, whereas a reduction in “contamination effects” represents a positive externality from private security);
- The game-theoretic nature of the interaction with terrorists, in the sense that the threat can respond to the security measure;

- The importance of “missing” markets, including imperfections in the reinsurance markets; and
- The relative emphasis to be placed on preventive activities (such as stronger border and visa systems) relative to protective activities (such as improved security at buildings) and consequence management activities (such as improved communications among first responders)

How can an NSF initiative most effectively promote future research in risk analysis and decision making that brings together the natural and social sciences?

One possibility is to focus initially on a specific problem, such as the protection of chemical facilities. Such facilities pose a potentially lethal threat to millions of people, yet little has been done thus far to protect the facilities against intentional sabotage (as opposed to more traditional safety concerns). Research issues could include the physical dangers associated with different types of chemicals; the ability of chemical facility managers and owners to evaluate the threat of potential attacks; the importance of the “contamination” effect from one chemical facility on other firms and population centers; and the role of insurance and third-party inspection services in improving security at the facilities.

In the Brookings volume, we argue that chemical facilities may be particularly amenable to a combination of regulatory standards and insurance requirements. A mixed system of minimum standards coupled with an insurance mandate not only can encourage actors to bolster security, but can provide incentives for innovation to reduce the costs of achieving any given level of security and can reduce the monitoring burden for government regulators. The presence of minimum regulatory standards also helps to attenuate the moral hazard effect from insurance, and can provide guidance to courts in determining negligence under the liability laws. Further research would be warranted to extend and expand the promising pilot projects conducted using insurance firms and third-party auditors to bolster compliance with the existing Section 112(r) regulations for chemical facilities.

Charles Plott

The outline below provides several different sources of scientific expertise that can be brought to focus on the problem risk analysis and decision making to address societal problems. The differences in approach are striking and should be reflected in any research agenda shaped by the NSF.

1. Single sources vs. multiple sources of risk

Risk can result from exposure to a single source or multiple sources of different types of risks – patterns of risk, so to speak. Much analysis is cast in terms of a single source, perhaps with a pattern of events simply labeled as a single event. However, if there are multiple sources of risk the analysis can benefit from the economics of uncertainty and related theories of portfolios of risk that teach us that modifications of one dimension of risk become undone by behavior that adjusts other dimensions of risk in response. As an illustration, (controversial) evidence has been advanced that suggests access to seat belts is accompanied by individual decisions to follow closer and drive faster. As another illustration, communication devices, more secure life preservers and boats that are not supposed to sink are accompanied by individual decisions to go farther offshore and in more dangerous weather. Such examples of “home made leverage” are well known from finance.

Thus, on the one hand portfolio adjustments may tend to frustrate or undo the effectiveness of some policies while on the other hand the same type of behavior can be seen as embracing and supporting other types of policy. Again, from the economics of uncertainty we have seen the discovery and implementation of imaginative instruments that help individuals remove risk from their portfolios. Options contracts, futures contracts as well as special types of mutual funds all are instruments and institutions designed specifically for risk coping needs. The discovery of these instruments, their designs and implementation rest on principles of the economics of uncertainty that could have much wider application as sources of policy tools than have been attempted thus far.

2. Single humans decision makers vs. other decision-making creatures

Many major decisions regarding risk are the result of a process that involves more than one person. Often major decisions involve more than one institution, each of which has its own internal decision processes. A fundamental lesson of social choice theory and public choice theory is that the principles that govern the decisions of groups of individuals are not the same as the principles that govern the decisions of individuals in the group. Institutions and decision-making rules play a decisive role.

Clearly a better understanding of individual decisions is important but any investigation of the individuals should be accompanied by research on the social decision, social choice and public choice processes. We know a great deal about the role of institutions such as voting processes in

conflict resolution and preference aggregation. For example, group decisions exhibit inconsistencies and inefficiencies even when each person in the group behaves rationally and consistently. We know much less about the role of those same processes when information is variable and belief aggregation becomes an aspect of the process.

There are also many more types of processes and institutions that deal with information and risk taking, almost none of which have been subjected to the same depth of analysis as those related to conflict resolution. These include processes that separate value determination, probability determination and decision into three or more different processes. Included also are processes or organization that involve the delegation of decisions, systems of advisors (multiple principles and multiple agents), aggregation of beliefs, or information processing institutions. In all such processes there exist the possibility of conflict of interest, information screening and other phenomena that might result in poor decisions. The theory is rapidly developing insights about institutions that work “properly” in spite of such problems.

The problems associated with the aggregation of preferences and rules of voting are well researched. But, even here aspects of time dynamics, irreversible decisions and changing or evolving event spaces (unanticipated phenomena) are not well understood.

3. Sources and types of information

Frequently when we consider risk analysis we have in mind nature facts, the results of experiments or statistical samples. However, much of our feelings of likelihood and the decisions we make regarding risk do not come from such sources. They come from observing the decisions of others. People gain information through observing the behavior of others. This is not mindless herd-like or mimicking behavior. It is a thoughtful process and is often an efficient method of information accumulation and aggregation.

The role of markets in transmitting information from those who have to those who do not has long been recognized. Market reactions are known to anticipate things to come as those with good information attempt to take advantage of those who do not and in the process change the price of the stock. Price changes broadcast what is known to all who are watching the stock movements. The phenomenon of people learning from the actions of others takes place in many forms, not only markets. If you and your friend stand in front of a shop window looking, others will pause and look. In this case and in the case of markets, the outsider is inferring something about the information held by another by observing decision behavior. The phenomenon is not unique to humans as is known to any fisherman who watches the behavior of birds or watches birds watch other birds. Indeed, much of what we know is acquired this way.

Not only does our understanding of this belief transfer mechanism have implications for broad policies, it has implications for the design of new types of information aggregation mechanisms. Recent research has been successful employing the principles to design special institutions that have only a purpose of capturing information that is distributed in small “amounts” across many people and is held in the form of subjective beliefs, impulses and feelings. Once captured and aggregated, it can be transformed from the “soft” form into a harder, more useful form of probabilities. Experimental work with such processes has been very successful and might have

application in the area of rare events, warning systems and other situations in which relevant data are dispersed across the experiences of many, isolated individuals. In some instances the probability is low for any particular individual so most individuals have no experience at all. However, in a large population there may be many such experiences that when aggregated are important sources of information.

4. Administrative decisions vs. the design of processes and institutions

On the one hand the research can focus on single administrative acts, a balancing of specific risks and benefits relative to a specific source of risk. An example would be a decision about whether or not a certain drug should be used. Or the research could focus on the behavior of decision-making systems in response to risks. When the focus is on decision-making systems, the role of information and incentives related to information and actions become central. The two examples that follow are illustrative of how systemic analysis can lead to much different perspectives.

Example 1. Faced with a threat of an individually directed terrorist act, each agent can at a cost, invest in erecting deterrents that cannot prevent the terror act but make it more costly. The effect is to shift the threat to the agent's neighbors who have not invested. The equilibrium outcome is that all individuals invest. Each individual is forced to invest up to the full economically feasible level of deterrence because those who have invested the least will have the threat focused on them. That is, the system adjusts such that the probability of the terrorist attack goes up for those who have not invested in the deterrence. It is not hard to produce examples in which the equilibrium result is that the threat to a given individual is the same as it would be if no one had invested. In the end, the investment in deterrence was a pure social waste. That is, the individual responses, each of which were rational and well calculated, led to no improvement in the risk exposure even though the cost to the system might have been substantial.

Example 2. Suppose each individual is connect to another through a network. Suppose further that if one individual "fails/gets infected/overloads/experiences error/bankrupts" there is a high probability that the neighbors will have the same experience. Consider a case in which for a small investment an individual can prevent his system from exporting problems but such investment does nothing to protect the individual from importing problems. Because of the nature of the externality such investments are not made and the system performance itself is at risk as are any activities that depend upon it.

As another case suppose the agents in the network are engaging in risky behavior but the activities of one are necessary inputs for the activities of neighbors. Each produces outputs that are used as inputs by neighbors. Assume first that while the increased intensity of output production involves increased profit for the agent it is also associated with an increased probability of the agent's system failure and resulting in no output at all. Thus, each producer pushes the "local" system to a proper balancing of risks and profit. However, if neighbors are unaware of the risks taken by those on whom they depend there is an uninsured and unprotected exposure to the risk of catastrophic system failure. In this case there is a natural tendency of network development to push itself until it fails completely.

The point of these examples is that systemic responses to information and risk are the product of organization and incentives. The methods of coping with such risks must be recognized and designed into the fabric of the institution.

Appropriate principles, such as incentive compatibility, and methodologies of institutional design and testbedding lend themselves naturally to the analysis and to applications of this form.

The observation that catastrophe bonds might be an instrument that supplies the need for risk sharing of low probability, very high consequence events is a good example of the type of policies that one can imagine when using the approach. Modern technologies, including rapid communication and computation capacities, are being used to design institutions, instruments and processes that were never possible before. Suggested new institutions such can be studied more completely through the use of experimental testbeds, thereby lowering the cost and facilitating the consensus needed for successful policy implementation.

David Schkade


Which Type of Utility Defines a Risk?

Evaluating risky options requires evaluating outcomes that have not yet occurred. Decision and cost-benefit analyses usually assume that evaluations of an outcome made *ex ante* will match those made at the time the outcome is experienced (e.g., preference stationarity). A new and growing body of research shows that evaluations made before, during and after an event or change of circumstance are not always the same, and often differ systematically. This possibility raises the question of what the appropriate measure of utility should be for measuring risk. This section expands on this theme and explores possible implications for research and policy.

Varieties of Utility


In the 1990s Kahneman introduced a new and richer conception of utility which mapped out several possible aspects, most of which are not necessarily reflected in choice. *Decision utility* is the modern concept of utility; it is inferred from choices and is used in turn to explain choices. *Experienced utility* is Bentham's (1789) concept; it is a measure of subjective hedonic experience of outcomes or events. Reports of hedonic experience can be obtained in real time, or retrospectively. *Remembered utility* is the label for retrospective judgments of particular experiences. Finally, *predicted utility* refers to forecasts of hedonic experience (what might be called the "lay theory" of utility).

There is already much evidence that outcomes are often evaluated differently from these different perspectives. One stream of work shows that experienced utility does not always follow predicted utility or decision utility. For example, when evaluating the impact of possible changes in circumstances people often fail to appreciate the role of adaptation. In a well known article in the psychological literature on well-being, Brickman, Coates, and Janof-Bulman (1978) reported that after one year there were only small differences in life satisfaction between paraplegics and normal controls on the one hand, and between lottery winners and normal controls on the other. A more recent study reported that the frequency of positive affect returns to normal levels within one to two years following the death of a loved one (Suh, Diener, and Fujita, 1996; see Figure below). Gilbert et al. (1998) studied assistant professors' forecasts of how they would feel at various points in time after their tenure decision, and compared these forecasts to the self-reported well being of others whose tenure decision had been made in the past. Current assistants predicted that they would be much happier during the first five years after a positive decision but, surprisingly, there was no significant difference in reported well being between those who had and had not received tenure in either the first five or the next five years afterward. People in California are no happier than people elsewhere, despite a wide spread belief that life must be better there (Schkade and Kahneman, 1998). Loewenstein and Schkade (1999) review these and many other examples of errors in forecasting future utility.



Another stream of work has found discrepancies between evaluations of experiences as they happen (experienced utility) and in evaluations of the same experiences in retrospect (remembered utility). Kahneman and colleagues (1993; 1996) proposed a “snapshot model” in which an experience is represented selectively in memory by a few prototypical moments rather than by the full experience. Retrospective evaluations are then performed on this partial representation of the experience. The central finding of these studies was that retrospective evaluations of experiences were predicted with substantial accuracy by a simple average of the Peak (the most extreme moment) and the End (the last moment). This has been called the Peak/End rule. On the other hand, the correlation between retrospective evaluations and the duration of the experience was negligible, a finding that Fredrickson and Kahneman (1993) labeled *duration neglect*.

Redelmeier and Kahneman (1996) presented a compelling example of the Peak/End rule and duration neglect in a medical context. Patients undergoing a colonoscopy reported their current level of pain every 60 seconds throughout the procedure. Later, subjects were asked to evaluate their experience overall, and the Peak and End pain predicted these retrospective evaluations. Further, while the duration of the procedure varied considerably among patients (from 4 to 69 minutes), it had no influence on retrospective evaluations. The implications of these psychological rules of evaluation are paradoxical, and somewhat startling. In the Figure, for example, it appears evident that patient B had a worse colonoscopy than patient A. However, it is also apparent that the Peak/End average was worse for patient A, whose procedure ended at a moment of relatively intense pain. In fact, as predicted from the Peak/End rule, it was patient A who evaluated the procedure most negatively, and was least likely to agree to a repeat colonoscopy.



Finally, there is the question of the extent to which standard patterns of preference in choice (decision utility) such as loss aversion also manifest themselves in other aspects of utility. Loss aversion is said to occur when a loss of a given magnitude has more influence on choices than does a gain of the same magnitude. There is preliminary evidence (Schkade and Kahneman, 2002) that the same gain-loss asymmetry does not occur in predicted, experienced and remembered utilities. While tentative at this point, in predicted and experienced utilities, the evidence points toward no strong pattern of negative events having a larger influence on evaluations than positive events. More broadly there is no evidence that bad dominates good in the memories of events; if anything, there appears to be a bias toward the good (Matlin and Stang, 1978). This observation suggests that there may be no negativity bias in remembered utility – and perhaps that there is none, more generally, in other aspects of utility. All this points to the possibility that loss aversion may not reflect an accurate anticipation of the experienced utility of gains and losses, and is it best understood as a cautious strategy of choice, i.e., as a feature of decision utility? If so, the discrepancy suggests that it is in some sense it may be a mistake.

Are We Optimizing the Right Thing?

If we assume for a moment that evaluations before, during and after and experience do differ, what are the implications for research and policy? A serious normative question is raised about whether we should be optimizing decision utility (probably the most common approach

currently), experienced utility or remembered utility in our analyses of risk. Should we focus on giving people what they want (decision utility) or on making their lives better (improve experienced or remembered utility)? In some cases, these objectives are probably aligned, but in others, as we have seen, they diverge.

In some sense we are already struggling with this dilemma in some policy debates. Many academic discussions of comparative risk (e.g., cost per life saved) emphasize historical statistics about events that have already occurred, while perceived risk would seem to depend on a complex combination of what has been and what might be (whether it has been experienced or not). Issues of paternalism are raised when physicians or policy makers make decisions that they feel are in the best interest of their patient or constituents, but which those represented might not choose themselves. In the colonoscopy example above, whose perspective should rule – the person’s perception of what a colonoscopy *will* be like, *is* like, or *was* like? If the objective of the health care system is to increase repeat colonoscopies, then it is clear that the quality of the experience as it occurs should be subjugated (i.e., by adding a period of decreasing pain at the end of the procedure) to the construction of a better memory of the procedure that would encourage a return visit. Yet in the middle of the procedure, the patient would surely reject this proposition. Should we be more concerned about improving the experience of being a flood victim or with decreasing the chance that they will make choices that make them victims again in the future?

These are deep waters and I do not claim to see clearly through them. But it seems likely that descriptive research designed to map out the relationship among these varieties of utility, and normative analyses of related policy implications would add valuable tools to the analysis and management of societal risks.

Paul Slovic

Risk as Feeling and Risk as Analysis

How do we make decisions in the face of risk and uncertainty? How should we make such decisions? What is the proper role of risk analysis and other analytic tools in decision making? What do we know about these questions and what more do we need to know in the world post-September 11?

1. Risk as Feeling

Risk analysis is not the only rational technique for making decisions in situations of risk. Modern theories of risk perception and cognition inform us that, besides risk analysis (and its close relative, decision analysis), we have another mode of thinking that is essential for rational decisions in the face of danger. This is the experiential mode, which enabled us to survive during the long period of human evolution and remains the most natural and most common way to respond to threat, even in the modern world. Experiential thinking is intuitive, automatic, and fast. It relies on images and associations, linked by experience to emotions and affect (feelings that something is good or bad). This system represents risk as a gut feeling, telling us whether it is safe to walk down this dark street or drink this strange-smelling water.

Proponents of formal analysis, the “newcomer” on the risk management scene, tend to view “risk as feeling” as irrational. It is not. Sophisticated studies by neuroscientists such as Antonio Damasio and others have demonstrated that logical argument and analytic reasoning cannot be effective unless guided by emotion and affect. Rational decision making requires proper integration of both modes of thought. However, both systems have their biases and limitations. The challenge before us is to figure out how to minimize these limitations when we assess risks. Thus, when our feelings of fear move us to consider purchasing a handgun to protect against terrorists, our analytic selves should also heed the evidence showing that a gun fired in the home is 22 times more likely to harm oneself or a friend or family member than to harm an unknown, hostile intruder. Risk as feeling tends to overweight frightening consequences. Risk as analysis can give us perspective on the likelihood of such consequences. Conversely, when our analytic selves dominate (example: the study by Philip Morris in the Czech Republic demonstrating that smokers save the government money by dying quickly and young), we need to look to our experiential system to cross-check the wisdom or the implications of such analyses.

In some circumstances, risk as feeling may outperform risk as analysis. A case in point is a news story dated March 27, 2002 discussing the difficulty of screening 150,000 checked pieces of baggage at Los Angeles International Airport. The best analytic devices, utilizing x-rays, computers, and other modern tools, are slow and inaccurate. The solution – rely upon the noses of trained dogs.

Some species of trouble – such as terrorism – greatly strain the capacity of quantitative risk analysis. Our models of the hazard generating process are too crude to permit precise and accurate predictions of where, when, and how the next attacks might unfold. What is the role of

risk analysis when the stakes are high, the uncertainties are enormous, and time is precious? Is there a human equivalent of the dog's nose that can be put to good use in such circumstances, relying on instinctual processing of affective cues using brain mechanisms honed through evolution to enhance survival? What research is needed to train and test experiential risk analysis skills?

2. Risk Analysis vs. Decision Analysis

Certainly we need analytic thinking as well as intuitive thinking in today's dangerous world. But is risk analysis the right mode of analysis? I think not. Risk is a complex and controversial concept that typically has no direct implications for decision making. Assessing a risk as "high" does not necessarily mean we should act to reduce it. Similarly, assessing a risk as "low" does not mean we should ignore it. Risk management decisions depend on the balancing of options, benefits, and other costs—not just risk. In this sense, we need to look beyond measurements of something called "risk" to make effective risk management decisions. In particular, we may need to embed risk decisions more strongly in techniques for sound social and individual decision making.

The techniques of decision analysis has been developed to help managers and policy makers make complex decisions in the face of risk and uncertainty. A decision analysis approach to risk decision making has several advantages over the conventional practices of risk analysis. As outlined in Table 1, these advantages can be traced to the grounding of decision analysis methods in a specific framing, or social context, for decision making. Thus, whereas risk is the central concept in risk analysis, with the problems noted above, the structure of the decision problem is central to decision analysis, along with the relevant probabilities and values.

Second, whereas risk analysts tend to conceive of risk as real and objective, and deride risk perceptions as subjective and emotional, decision analysis respects the subjectivity of probabilities and outcome values. This shows up most clearly in the process by which those impacts designated as important are defined as a part of each individual risk assessment. For example, if fatalities matter, then the next question is "why?" Do they matter because humans or animal species are affected? Do they matter because of the number of expected fatalities or because of the emotion associated with particular modes of dying?

Third, risk analysis is distinctly "expert-centered" and uncomfortable (even hostile) toward considering the views of diverse, non-expert parties. It is also uncomfortable with a broadly multidimensional view of risk. In contrast, decision analysis seeks out the diverse views of interested and affected stakeholders. It attempts to assess the probabilities associated with all the outcomes believed to be important, and it assigns values to those outcomes in ways that can be sensitive to equity, personal control, catastrophic losses, or other factors deemed important by the affected parties.

Fourth, whereas risk analysis often strives for some magic number that defines an "acceptable risk," decision analysis recognizes that there is no universally acceptable level of risk. In decision analysis, acceptable risk depends upon the problem context and can be understood only in association with the management option that is best in a particular context. In other words, acceptable risk is decision driven: as the decision changes, so too will the magnitude of the risk that is acceptable (that is, the probabilities, consequences, etc. that are acceptable).

Table 1. Differences Between Risk Analysis and Decision Analysis

<i>Risk analysis</i>	<i>Decision analysis</i>
1. Risk is the central concept	1. Problem structure, probabilities, and values are central
2. Subjectivity is respected and incorporated into the analysis	2. Subjectivity is respected and incorporated into the analysis
3. Models the multidimensional views and values of interested and affected parties	3. Models the multidimensional views and values of interested and affected parties
4. Seeks acceptable level of risk as a standard to attain across problem domains	4. Acceptable risk is context-dependent and decision driven

The use of a decision analysis framework for addressing risk management problems allows the tools of modern risk analysis to be used as a part of a broader context, where the emphasis is on creating a sound structure for decision making rather than addressing the narrow concept of risk as some chance of loss. Because decision analysis avoids reification of the ill-defined concept “risk,” because it provides a more inclusive framework than risk analysis, and because it recognizes interested and affected parties as legitimate partners in the analysis, it deserves greater recognition and use in risk-management decisions.

3. Concluding Thoughts

The world has always been a dangerous place and today’s risks are certainly no less daunting than these of earlier eras. However the nature of these risks has definitely changed, posing unprecedented combinations of relatively high probability, catastrophic outcomes surrounded by vague knowledge and substantial uncertainties. Coping with such threats requires all the skills we can muster. This means tapping into the strengths of trained human judgment and intuition and integrating these with formal methods of analysis. Research is needed to show us how best to do this.

One additional thought. There are hundreds of researchers who have spent careers studying judgment and decision processes. Much is already known about the strengths and weaknesses of these processes. Procedures need to be installed to insure that the wisdom we already have is applied effectively to today’s challenging decision problems.

One “small” example:

The insurance industry is struggling to figure out how to assess probabilities of losses from various types of terrorist acts. Human judgment is essential for this task. How should it be applied? There is a research literature that is relevant to this question and there are researchers and practitioners knowledgeable about the application of this knowledge. They should be consulted.

Cass R. Sunstein

Debiasing? Regulation?

"The nation is quickly buying up stocks of gas masks, shelves are being stripped of antibiotics, and bottled water may not be far behind. Many travelers have canceled trips by air and taken trains or cars instead, even across the country. New Yorkers fearful of an attack on the subways insist on riding in cars on traffic-choked streets. Doctors in Boston report that patients with minor ailments like colds and sore throats have been calling out of fear that they may have been sickened by a toxic chemical or lethal germ introduced by terrorists. Meanwhile, business at McDonald's and Haagen-Dazs is thriving. What does this say about how people respond to potential threats to their health and lives?"²

It is widely agreed that regulatory policy suffers from a pattern of "paranoia and neglect" – a situation in which some risks receive a great deal of attention, even though they are statistically small or even tiny, and in which other risks receive little attention, even though they are statistically large, even massive. A glance at the "cost-per-life-saved" charts will be sufficient to make the point. There is also good reason to believe that at the individual level, people show a similar pattern, fearing some risks that are statistically small but offering little concern about some risks that are much larger. A central question is: What can government do about the problem (assuming, as I believe, that it is a problem)?

An obvious possibility, on which we now lack sufficient knowledge, involves the conditions under which information and disclosure will produce sensible behavioral changes. When will information campaigns work? When will they induce panic? When will they produce yawns? Are there demographic and cultural differences here? What are they?

We know that people are not sensitive to differences in low-probabilities (they do not show a materially greater reaction to a risk of 1/100,000 than to a risk of 1/1,000,000). We know that when emotions are engaged, people are especially insensitive to large variations in probability. (Hence "probability neglect" is pervasive.) We know, thanks to Paul Slovic and others, that people's reactions to risk are often influenced by their "affect" to the product, process, or activity in question – and that the affect heuristic can produce large blunders. We know too that people use the availability heuristic to assess risks, so that a salient example can have undeserved power (and the absence of a salient example can help produce risk neglect). We know that trust matters a great deal (consider the British fear of eating beef, produced in part by a loss of trust). We know finally that new and unfamiliar risks often receive excessive attention, and that old and familiar risks often receive too little attention. All of these effects can produce distortions in individual behavior and in law and policy.

What we do not know is whether and how government can counteract these various effects and make people appropriately sensitive to risk levels. Sometimes government knows that a risk is high or that a risk is low, and it would be very important to be able to inform people of the

² Jane E. Brody, Don't Lose Sight of Real, Everyday Risks, N.Y. Times, Oct. 9, 2001, at F6.

relevant facts. (It is worth thinking about some kind of “national warnings system,” if it could be made effective.) Sometimes government cannot quantify risks in any serious way, but it nonetheless needs to be able to prevent public panics when the evidence does not warrant them. The simple suggestion is that it would be desirable to produce “debiasing” strategies -- to know which ones work, and why.

An additional question is when and how debiasing strategies should be supplemented with regulatory/legal approaches. This is partly a normative question, not only an empirical one; but it requires a great deal of empirical work. Often social norms are extremely helpful in protecting against serious risks; sometimes social norms are actually harmful, ensuring that people will run serious risks. When can regulation start, or strengthen, desirable norms? Or reduce the hold of weak ones? There are hard questions here about the legitimate scope of paternalism. But a central point is that fear is itself “a cost,” and sometimes a high one. It can also induce large additional effects, some of them causing significant harm, including harm to health and safety. Hence government regulation that reduces fear might be justified even if the fear is baseless. What kind of regulation reduces fear? And is there a way to monetize the costs of fear itself, as well as its likely ancillary effects? It should be possible to make large advances on these questions.

Peter Szolovits

I suspect that many of the critical issues in developing a societal understanding of and response to risk center on politics, communication of appropriate metaphors to the public, and a spirit of dedication and vigilance on the part of the many who serve the public, from airline baggage checkers to postal workers to policemen to public health officers to policy formulators. Nevertheless, there are a number of technical research directions that may be helpful. I hope to learn much more about these topics from discussions with my colleagues at the workshop, but here are a few preliminary thoughts, which may be broader than the intended focus of the Workshop.

Broadening the Vision

As a very late contributor, I have had the benefit of being able to read the comments of most of my colleagues. They are striking not only for their methodologic and practical insights but because they show how much the field appears focused on serving the needs of policy makers and analysts, to the exclusion of those with line responsibilities either in anticipation or in response to disasters. The other notable bias is that risk analysis and its brethren receive far more attention than decision making, which I think of as the place where “the rubber hits the road.” (With the main exception of Slovic’s paper.)

My first, perhaps naïve, suggestion would be to make sure we take a broader view to include others. In my conception, decision makers fall into two major categories:

1. Those who set policy in a relatively static world, where there is time to do careful analyses, subject decisions to public and political processes, and fine-tune them based on experience, if any.
2. Those who are “on-line” in trying to anticipate, react to, rescue the victims of, and restore infrastructure in the wake of potential or actual disasters. This is probably a much larger community, many of whose members may not even know that they will play such roles until a disaster happens.

I would like to be sure that the agenda of our Workshop and the possible NSF program it leads to takes the needs of this second category of decision makers into account as well as the first.

Communications

One of the key common factors in responding to disasters, whether man-made or natural, is the inability of vast portions of society to communicate with one another. I presume that organizations such as FEMA have developed effective emergency communication tools within their own organizations and perhaps to some selected sister agencies. It still appears that to first approximation the kind of *ad hoc* communication whose need is recognized only during and after a disaster is poorly served by our normal communication infrastructure. For example, apparently in NYC on 9/11/2001, both wired and wireless phone service were severely disrupted. This

means that most people trying to cobble together appropriate responses to the disaster had no customary way of coordinating with others. Many research groups are addressing how to build communication devices that adapt automatically to whatever system is reachable around them, and there is much work to do in self-organizing networks, very robust directory services, etc., that could allow a degree of normal communication if any systems at all are functioning. On a more organizational level, until last year it was impossible for anyone to find out easily how many people were seen at each (or all) of the dozen or so emergency departments in Boston. To borrow a term from computer science, it would be immensely useful to develop a communication architecture that allows us to analyze the existing and needed flows of information both in anticipation of and in response to disasters, and to study the benefits and costs (e.g., to privacy, commercial concerns) of principles and practices that might be adopted. I believe that there is a set of (technical and social) engineering concerns of this type that are of immense practical importance.

Anticipation

Conventional wisdom in the biodefense community is that the number of fatalities from an attack rises as a sigmoid function of the time of detection, because delays in prevention or treatment of biological agents is devastating. Clearly, the best point of detection then is before an attack. This is normally the province of intelligence agencies, where indeed decades of work have been done toward comprehensive data collection, intelligent data fusion, etc. Nevertheless, this seems an interesting and potentially fruitful area for further research. Anecdotally, we hear of human analysts being able to spot and predict obscure trends and events. It may be worthwhile to undertake serious cognitive studies of how human experts can do this, and to develop a variety of approaches to augment their abilities. This may be in the form of “expert systems,” data mining methods, or more classical statistical or Bayesian approaches.

Modeling

I have seen many very interesting suggestions by others on topics in methodology and modeling. I would suggest only that we might want to add some other types of models. First, unlike risk analysis for natural disasters, where, with Camus, we may count on the indifference of the universe, terrorist acts are deliberate and planned by intelligent adversaries, so we cannot use expected values but should be thinking about the mini-max models typical of game theory. This is recognized in examples that point out the “weakest picket in the fence” problem, but probably deserves more formal adoption. Second, many of the on-line decisions in dealing with disaster must be made in real time. Therefore, approaches that take into consideration the cost of delays imposed by the decision-making process itself are critically needed. More generally, decision making involves building a decision model that includes possible events, their likelihoods, elaborating the set of possible actions, assessing the likelihoods of their outcomes, assigning value to ultimate outcomes, etc. Each of these steps can be brought within the overall decision-making model, so that the meta-level decisions that determine how to make a decision are themselves subject to analysis. Third, to the extent that disaster scenarios may be formally modeled, methods such as hidden Markov processes and partially observable Markov decision processes may provide an appropriate language. This is the focus of much current robotics research, where the formalization of how a robot is to make decisions under uncertainty is truly central.

Kathleen Tierney

Introduction

Our society and culture will be dominated by concerns about terrorism and weapons of mass destruction for years or not decades to come. The sudden emergence of these concerns after September 11 raises questions, not only for government officials and the public at large, but also for members of the scientific community whose work centers on risk assessment, perception, communication, and management. NSF clearly has a very significant role to play in developing and supporting a research agenda that includes both fundamental and problem-focused research on a range of themes related to homeland security. This paper will outline a series of general and more specific research questions that should be addressed in the new program NSF is developing. Some of these questions focus on the extent to which knowledge and practice derived from research on better-known hazards can inform the societal and institutional response to newly-emergent threats. Clearly, in whatever research is planned, it will be important to build upon what is already known and to focus on comparisons and contrasts, continuities and discontinuities between previously-identified hazards and those that have recently emerged.

What Problem Areas and Topics Should Be Considered?

1) Risk Perception, Communication, and Public Action

Comparing Hazards

New research should seek to determine the extent to which research on chronic and acute hazards is transferable to new homeland security threats. Chronic hazards include ongoing health, environmental, and social threats (e.g., crime). Acute hazards encompass rapid-onset natural and technological disasters. Based on decades of research, there is a good foundation of knowledge concerning how the general public (and to a much lesser extent, different segments of the public, such as minority group members—see discussion below) perceives various risks, how individuals and groups respond in the face of hazards, what factors influence perceptions and actions with respect to hazards, and what measures appear to work in encouraging appropriate behaviors, such as the adoption of self-protective measures. What is not known is the extent to which this knowledge applies to newly-emergent, highly diffuse, and dread risks such as terrorism and bioterrorism. Indeed, at this point, we have little grasp on what our highly diverse citizenry actually does know about these threats—a fundamental question that warrants immediate study.

What we do know suggests that homeland security threats present special and perhaps unprecedented challenges with respect to risk communication, warning, and, at a more general level, overall management. For example, in the field of disaster research, various hazards (or disaster agents) are typically characterized along a number of different dimensions that affect how they are perceived and managed. These include familiarity, understanding of, and experience with the hazard; short- and long-term predictability; speed of onset; duration of onset; extent to which the hazard is accompanied by perceptible physical precursors; scope of impact;

potential for repetitive, as opposed to one-time impact; and the extent to which the hazard is a consequence of human action or negligence, as opposed to “natural” causes. Other work on risk perception has identified other hazard parameters that also affect risk perception and action, including familiarity, dread, and the potential for multigenerational and inequitable impacts. Seen from this perspective, a hazard such as bioterrorism is unfamiliar and poorly understood to the vast majority of the public; it is outside the range of normal experience; its dissemination cannot be predicted; and there are no identifiable prior warning signs. There is the potential for spread of the hazard over a large but indeterminate area, multiple attacks are possible, the public can be expected to react to news of an attack with a very high level of concern, and at least in some cases there is the potential for multigenerational effects from exposure. Moreover, unlike natural disasters, there is human blame involved in terrorist attacks. The research challenge is to understand how these hazard characteristics influence perception and behavior.

Societal Diversity and Hazard Response

As indicated above, there is a substantial body of research on the manner in which the US public responds to various types of hazards and hazard advisories, from generalized risk information to more specific warnings of imminent disaster impact. However, despite the fact that our society has grown increasingly diverse, we lack a more nuanced understanding of hazard perception and response among minority and ethnic populations in the US. Our many minority populations differ from the Caucasian majority along a variety of dimensions that are known to be associated with hazard-related perceptions and behaviors, including language, income and educational levels, knowledge of hazards and hazard adjustments, trust in science and in governmental institutions, patterns of media use, and family size and composition. However, with a few notable exceptions, almost no systematic research has been conducted focusing specifically on minority populations and hazards. What is known about these populations suggests that their knowledge and actions with respect to hazards may well vary markedly from those of the majority population.

Particularly in the last few years, many hazards scholars have pointed to our lack of understanding of the ways in which societal diversity influences knowledge and behavior with respect to hazards, and there have been many calls for additional research. This knowledge gap now appears even more important in light of the new homeland security threats our society faces. Although it is conceivable that the next terrorist attack will take place in a suburban or rural setting, it seems much more likely that our major urban centers will be targeted. For example, it would make little sense to disperse a biohazard outside a densely-populated area. Given the tremendous diversity of our major urban centers—the very areas that terrorists are most likely to target—much more research is needed on minority community responses to hazards.

2) Institutions, Hazards, and Risk

Mass Media and New Communications Technologies

It is impossible to overstate the role of the mass media in shaping public perception and response to hazards and risks. While many in our society are exposed to hazard-related topics through the educational system and through other means, such as social network contacts, the vast majority

of what we know either comes directly from or is filtered through the media. This is particularly true during crisis situations, when public information-seeking intensifies and people turn to the media for information and advice on what actions they should take.

In less than one generation, there have been enormous changes in the mass media sector in the US, with greater complexity, diversity, and audience targeting on the one hand, and greater consolidation and centralization of mass media organizations on the other. With these large-scale changes, people are getting information from more sources than ever before, even as media content has in many cases become more homogeneous. In thinking about hazards and risk, this raises many questions, examples of which include: To what sources do different groups within the population turn for information on hazards, and what do they understand about hazards as a consequence of their media use? What types of media information are people more likely to retain, and how is that information ultimately used—if it is used at all? What media sources are seen as credible by different groups within the population?

In a related vein, there is also a need for more research on processes involved in news production and media framing of risks, with a specific emphasis on newly-identified threats. How do news organizations cover hazards, and what influences their selection of topics on which to report? How adequate is media reporting on hazards—particularly emerging homeland security threats? What “expert” information sources do different media use? How do the media frame these threats, and for what different audiences? Which media tend to be the most accurate and complete in their reporting? Which disseminate inaccurate information and myths—and which segments of the public rely on such sources for threat-related information?

A third area in which research is needed centers on the agenda-setting function of the media, both with respect to better-known hazards and with respect to homeland security threats. There is a need to explore linkages that exist among the media, general public opinion, positions taken by stakeholders and opinion leaders, and the construction of problems and issues within public policy arenas.

Among new communications media, the Internet stands out as a major focus for future research. As a medium to which tens of millions can simultaneously turn for information about hazards and risks, the Internet has the potential for disseminating massive amounts of information on a near-real-time basis. For example, through new Internet applications, the US Geological Survey has recently had very significant success in providing information on earthquake impacts when they occur and in obtaining information from residents of affected regions about their earthquake experiences. However, addition to transmitting sound and accurate information, the Internet also serves as a major source of misinformation and adds additional noise to already complicated communication processes. Again, this appears to be particularly true with respect to the threats with which the public has been preoccupied in the wake of 9-11. At worst, the web has become a medium for the promulgation of all sorts of conspiracy theories and misinformation regarding the terrorist threat. In light of its importance and potential societal impact, new research is needed on the role of the Internet in communicating risk-related information. Equally important, there is a need to better understand the role of the web as a medium facilitating the emergence of hazard-focused knowledge and action networks.

Trust in Institutions

Finally, recent events point to the need to explore in more depth questions related to public trust in the institutions that are charged with managing risk. Since 9-11, many things have taken place that have the potential for eroding public trust and support for institutions, including reported failures on the part of key intelligence and law enforcement agencies to identify early signs of an impending terrorist attack and what is widely seen as a mishandling of public communications and warnings regarding bioterrorism and other terrorist threats. The Enron debacle, revelations about stock analysts and accounting firms, and the regular release of new information on ways in which major corporations have used deceptive practices in reporting financial data also serve to undermine public trust. These recent events and problems suggest many researchable questions: To what extent have patterns of public trust been altered by events such as 9-11 and crises involving US financial institutions? Which institutions do different sectors of the public trust, both as information sources and as providers of solutions for societal risks? What institutions and organizations are seen as trustworthy sources of guidance on newly-emerging threats? How are trust and credibility sustained in situations such as those involving terrorist threats, in which false alarms and the “cry wolf” phenomenon are virtually inevitable? How do trust and recreancy affect perception and behavior with respect to risks and hazards, including in particular newly-emerging threats? If trust is lost, how can it be restored?

How Can an NSF Initiative Promote Future Research in Risk Analysis and Decision Making that Brings Together the Natural and Social Sciences?

I have no specific suggestions at this time. No doubt many useful ideas will emerge from the workshop itself.

Richard Zeckhauser

What follows are thoughts on three different subjects.

I. Change in the nature of risks of concern

The risk literature has most commonly discussed unintended risks, and distinguished those that are externally created from those a decision maker (DM) brings on herself. These risks fall into many, at times overlapping categories, risks created by nature, self-imposed risks (e.g., cardiovascular disease), and man-made risks involving others. The last, which are often the byproducts of desirable activities and treated as externalities, may be qualitatively different if they operate at a micro scale (e.g., auto collisions) as opposed to emerging from an aggregation process (e.g., stock market bubbles).

The two most salient news events of the past year, one in international affairs and the other relating to the domestic economy, involve risks of a quite different nature. Both are man-made, externally generated, and intentional (purposefully created). They are the September 11 terrorist attack, and the massive misreporting on earnings and expenditures by some prominent corporations. The former is a malicious risk. It entails costs to the perpetrator to bring harm on the victim, implying the risk creator values positively the victim's harm. The latter enables the perpetrator to usurp value from the sucker, though presumably she would prefer not to hurt him. Given the human intentionality behind these risks, any reasonable calculation would dramatically raise the estimate of the future likelihood of such events absent major policy measures. Even with such measures, the estimated risks may escalate despite the fact that we will be substantially more vigilant against them, will attempt to tear them out by the roots, and will erect significant deterrents against them.

Uncertainty, the ability to surprise the other party, is a major element of these intentional risks. Otherwise, the nations that were subject to terrorism could readily defeat their weak attackers. The strength of the victims is the source of their vulnerability. They possess simply too many valuable targets to be fully protected; given their wealth, they value life too highly to slough off what may be a small statistical risk of fatality.

Similarly with the corporate shenanigans. If expected, no one would invest in the companies of the perpetrators. Hence, the imposed costs must be difficult to discover. Many only become significant when unforeseen unfavorable external events occur.

Treatment of four risk classes. Humans may experience risks differently, either in utility terms, willingness to take gambles, or both, depending on how they arose. How will DMs treat risks differently depending on whether they are (1) malicious, (2) usurping or sucker-oriented, (3) external unintended, or (4) self-created? For example, will their assessments of subjective probabilities relative to realized probabilities vary over the four cases? Will a major component of the outcome depend on the source of the risk? For example, will people pay "too much" to avoid being a sucker?

FOUR CLASSES OF RISK				
Category	Unintended Byproduct Risks		Intentional Risks	
		Self-Created	Created by Others	Value Usurping (Sucker)
Example	Heart disease	Auto collision Pollution-related health loss	Corporate Earnings Manipulation	Terrorism

Misordered and misdirected responses. Particularly in approaching new risks, we are likely to significantly misorder our priorities. Such misordering is a well-established problem in dealing with risks, with big ones slighted and small ones tackled at great expense. Much misdirected effort is "barn door closing". The proposed nationwide protection of airports after the recent LAX shooting provides an example. There are literally tens of thousands of unprotected locations nationwide where a terrorist (or crazy guy) with a gun could kill more people than in an airport. Even recognizing the symbolism of airports, and El Al counters, such proposals make little sense. (Note also the proposed redesign of the recently announced \$6 billion plus reconstruction plan for LAX.) In general, it must be recognized that the value of protecting one location is only that the terrorist is driven to his next favorite, and often only slightly lower-valued location.

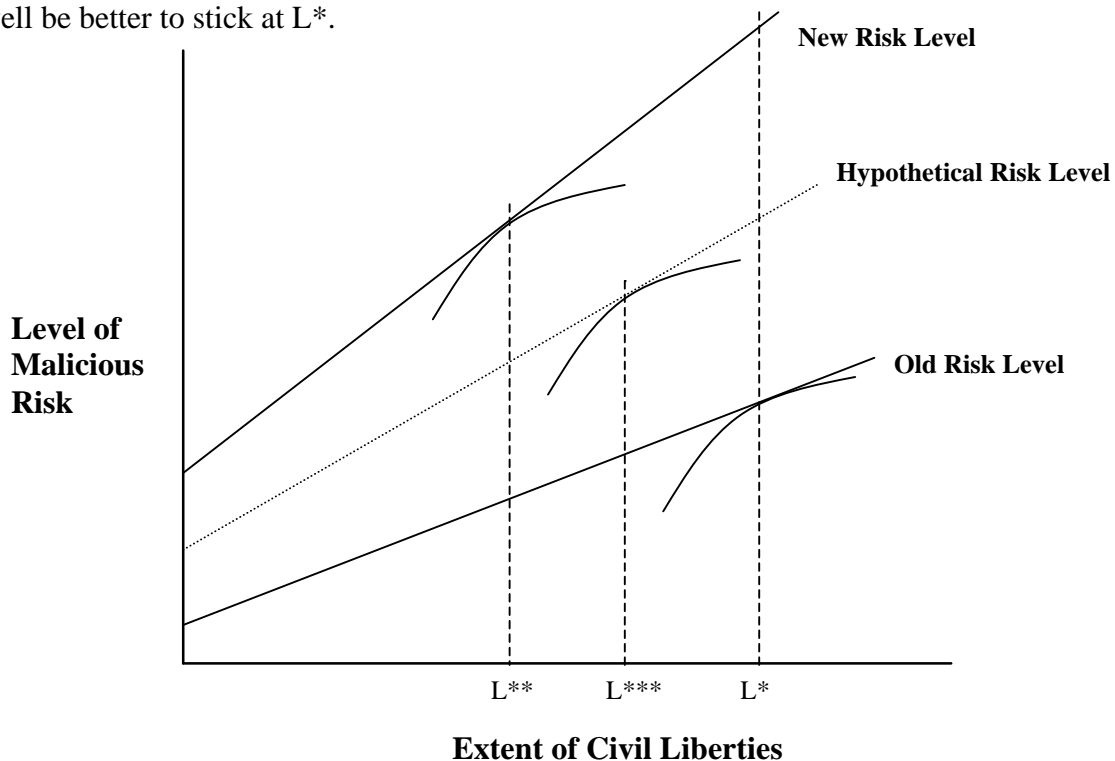
The corporate earnings scandal suggests a different class of problem. In the political heat of the moment, we are likely to rush to legislation, with little understanding of the scale of the problem, the ability of the government to assume tasks it will take on itself, etc. The general goal is to restore faith in the corporate governance system. Does this require deterring many small distortions, or a few big frauds, or possibly quite a few big frauds? There are likely to be strong tradeoffs, and the appropriate policy measures would differ depending on the mix of the problem and which parts are attacked. Are accounting firms no longer to be trusted, or does Arthur Andersen's speedy demise make all of them wish to be pure and true? The challenge here may be to escape from our current bind, where many corporations have distorted or misreported, but have no easy way to "get off the tiger". This class of disclosure problems extends, incidentally, to a range of risk issues, such as the health risks associated with various products that corporations produce.

Some promising policy proposals have surprising aspects. Most current measures being considered represent strong sticks. When an unexpected bad situation is discovered, it may be the best time to introduce some carrots as well, however politically unpopular that may be. Thus, we may have legislation requiring corporations to fess up to past misdeeds. (The vigorous prosecution of fraud is clear; these are thoughts for gray areas.) However, many past perpetrators will risk a further cover-up, and its greater penalties, rather than reveal. A promise of relief from certain penalties if confession is made by a certain date, may make revelation attractive, without too much sacrifice. Moreover, if many fess up, it will provide an informational externality about possible sins by others in like circumstances. Even the suspicion

that others will confess will spur information provision. -- The principal goal of policy should be to get the information out fully, lest financial markets fail to regain their confidence. (The decade-long experience in Japan with presumed insolvent banks is instructive.)

New risks and the sacrifice of old benefits. When new risks spring upon us, it may be worthwhile to employ dramatically new approaches, and to relax old and sacrosanct constraints. The situation was not foreseeable at the time the constraint was put in place. To illustrate, the figure shows the tradeoff between physical risk -- the negative of safety -- and civil liberties. In the old days, there was a firm constraint for civil liberties at L^* . With the new risk, say the terror attack, upon us, it is not merely that the safety level has increased, but its tradeoff against L has become steeper. The new optimal level is L^{**} .

With less dramatic escalations in risk, sacrifices in L would not be justified. Due to loss aversion, the cost would be too high. Thus, if the safety-civil liberties curve had shifted to the hypothetical line, even though the optimum assuming no history would have been at L^{***} , it may well be better to stick at L^* .



II. Amelioration, avoidance, and abatement

When economists think of risks, familiar marginal benefit and marginal costs leap into their heads, and they search for the crossing point. The control of some environmental pollutant might be their analytic metaphor. Many risk problems do not fit this neat model. For many risks, there are multiple, quite disparate ways of reducing the impact. Thus, the authorities can ameliorate the risks of harm in a terror attack by keeping appropriate medical supplies on hand, or arming sky marshals. An individual can reduce her risk through avoidance, e.g., by staying away from attractive targets, such as crowded places. Finally, the authorities can reduce or abate the risk itself by seeking to destroy the terror networks.

When considering their amelioration and abatement strategies, the authorities should consider how this will affect avoidance actions by citizens. If terrorists search for the most attractive targets, citizens' avoidance actions will impose significant externalities on each other as they seek to reduce their own risks.

Consider an extreme example for a homogeneous population. Per capita risk at a public event -- e.g., July 4 fireworks -- depends on the numbers attending (greater numbers attract terrorists) and government-provided security. Given any level of security, individuals will divide in equilibrium fashion between watching at home at attending. As long as there remain people watching at home, there is no benefit to increased security.³

Abatement strategies will frequently offer economies of scale. If so, the optimal strategy is likely to be continual amelioration, but at variable levels of intensity, with occasional massive abatement actions. We employ such strategies to clean up major beaches (where we occasionally replace a layer of sand), to confront life-threatening illnesses (continual lifestyle prevention with massive interventions after say a heart attack), and to deal with terror threats.

III. Policy poker: information games with partisan players

High-stakes decisions are often made through a political, legal, or bargaining process. This is the case in domestic policy in the United States and elsewhere, and in international settings as well. The process may be highly formalized, with a court or arbitrator adjudicating the results. More commonly it is a debate, carried out in scientific, bureaucratic, legislative, and diplomatic forums and in the media, with the public an important onlooker. But in virtually all cases, it is partisan: contending factions have different views of what should be done. Winning the policy debate usually requires winning the hearts and the minds of those without strong initial views among the public, officials, or legislators. Such is the way that the Vietnam War ended. Conversely, the current lack of a deep public concern about global climate change has surely diminished U.S. policy efforts in that area.

Our core argument is that partisan or adversarial decision processes introduce key biases in the way information is generated, revealed, and interpreted. This in turn, dramatically affects policy outcomes. At the workshop, we hope to stimulate a discussion of such biases. Here we mention but a few.

In high-stakes policy disputes, contending factions rarely pursue their preferred policy positions simply by issuing policy statements, or even by setting forth promises, threats, and demands. When stakes are high and debates are at least partly public, getting what you want usually requires persuading not just your friends and allies, but often those who are initially neutral, or even some who are opposed, that what you propose is right, prudent, or proper. As Elster has pointed out, this requires advancing arguments that at least appear to be impartial. They come in two forms: (1) based on normative principles, or (2) drawing on the implications of objective evidence for desirable action.

³ Essentially this is an equilibrium in one of Schelling's Micromotives and Macrobehavior situations, with one of the payoff curves being flat.

We are concerned with the second type of argument, which translates the implications of objective evidence into preferred actions. The results often depend on the translator. We illustrate for a range of policy settings: Thus, partisan factions argue over whether evidence of an economic slowdown justifies a drop in interest rates or fiscal stimulus; whether evidence of clandestine weapons development programs in Iraq merit an attack; whether genocidal activities underway in Rwanda, call for a response (whether military or diplomatic); whether the pattern of accidents and deaths involving Ford Explorers and Firestone tires justifies a recall; and whether the sharp increase in recorded temperatures in the late 20th century justifies restrictions on emissions of greenhouse gases. Sometimes these arguments are based on private information, but more frequently – because it is more persuasive -- the supporting evidence is common knowledge. The interpretation of public information as predictive evidence, however, is likely to be strongly contested. Even more strongly contested will be its ultimate implications for action.

Policy poker. Much relevant information is private, and may be selectively revealed or suppressed depending on whether it is likely to support or harm the preferred position of the party who holds it. In theory, if it cost nothing to reveal such information and everyone knew that the party held the information, there would be a complete unraveling, and all private information would become public. By this argument, we would expect pharmaceutical manufacturers to reveal all complaints about side effects of products on the market, since we know they hold this information and the mechanics of revealing it would cost little. In practice, many knowledgeable parties, including drug companies in this case, conceal even information that others know they possess. Presumably they are capitalizing on biases, a public that is more "out-of-sight, out of mind," as opposed to a master card player.

More generally, the selective commissioning of studies that are expected to lead to certain results, dissemination of information, framing of the information that is disseminated is all part of a game to influence the beliefs of other parties, and ultimately policy outcomes. This information-revelation process is all part of a game we have labeled Policy Poker.

This game, like many parlor games and poker itself, is rarely played at the highest levels. The skilled player knows how to exploit the biases of adversaries and onlooker participants. Collective processes often either exacerbate individual biases, and/or create new biases of their own.

Status quo (SQ) bias and anchoring. It is well known that individuals tend to stick with decisions they have already made, even if they had only minimal basis for the original decision. Similarly, they tend to anchor on original assessments and estimates, and to change them only sluggishly. This bias carries over to group processes, but is reinforced. In part, this is because individual SQ bias carries over, but group processes reinforce this tendency. Just as humans do not like to incur the costs of changing their minds, groups put processes in place so outcomes do not shift too swiftly. Thus, courts adhere to precedent in their decisions, legislatures prevent the same vote from reappearing, and boxing championships are taken away only when the loss is decisive.

When it comes to policy, those having differential sway over a policy decision tended to get their way, which reinforces their influence. This makes it quite possible that the group that has substantial influence will be difficult to dislodge, and like heavyweight boxing champions, will

only find its preferred policies supplanted when strong evidence come in for the other side, or where some other salient group rises in opposition. The beneficiary of this "winners' advantage" approach may come from any of a variety of constituencies. In the environmental examples we shall be presenting below, the groups in opposition to the environment were initially ascendant, but lost position once evidence grew up on the magnitude of risks. However, with risk regulation for food substances containing carcinogens (Delaney Clause)⁴ and FDA drug approvals, the consumer protection activist community was disproportionately influential, but subsequently lost ground. There is strong anecdotal evidence that the sea change in swifter time-to-market drug regulation is significantly due to pressures from the AIDS-activist community.

Suppression of contrary information. Individual decision makers, through forces related to cognitive dissonance and the like, suppress information they themselves possess, which is contrary to other beliefs they hold, or positions they have taken or intend to take. In group situations, this tendency may be strongly reinforced. It may be extremely costly for anyone to step outside the group, and make even mild statements, such as: "The evidence is not as strong for our side as others are saying." Such group pressures even operate when the group is meeting solely with supporters. Such a process for conformity in views was strongly in evidence this fall with the Florida aftermath of the Bush-Gore election. Neither on television nor in academic parlors, did we hear much discussion by the supporters of one candidate or the other of factual information that might be favorable to the other side.

Hyperbolic discounting. As Laibson and others have documented at length, individuals tend to discount the future excessively relative to next period, while not applying equivalent discount rates between successive periods in the future. The result is a time-inconsistency problem. Such hyperbolic preferences severely compromise situations where there is a tradeoff between current and future benefits, or current and future costs. Current benefits and current costs get over-weighted.

This process is reinforced where groups, not individuals, are making decisions, particularly if the action requires some reversal from the status quo. First, it is detrimental to the leader, and may cost his head, to be the first to suggest a change in policy. Second, the leader will often be in power for only a limited period of time. If he can avoid making the change "on his watch," the problem will fall to others. Thus, no cigarette company executive in the 1980s had much incentive to change the fortress attitude that the companies put forward relative to the dangers of smoking.

Excess attention to salient events. The public does not attend to subtle signals and indicators of the consequences of actions, even those that have strong predictive tendencies. And when it does, it gets confused because there are always authoritative voices arguing that it is not really significant. Thus, it takes a salient event, such as Three Mile Island, the burning of Yellowstone

⁴ The Delaney Clause, enacted in 1950, forbade the marketing of canned or processed foods with any detectable carcinogen residues. As the ability to detect chemicals advanced from parts per thousand to parts per billion over the subsequent five decades, food processors increasingly argued that meeting the requirement of the Delaney Clause was impossible. The clause was finally eliminated in the 1996 Food Quality Protection Act, with the acquiescence of environmental groups who in return gained agreement to conduct a comprehensive review of all pesticides to determine safe residue levels to protect children's health.

Park, or the Thalidomide tragedies and subsequent ban in the United States, to get public attention. The first event closed down nuclear plant construction in the United States ever since, despite great subsequent improvements in the economy, reliability, and safety of operations at existing plants. The heat wave and drought of 1988, highlighted by the burning of Yellowstone Park, brought global warming onto the policy agenda, first in the United States and quickly thereafter in Europe, even though that summer was uncommonly cool in Europe. The tragedy averted in the U.S. by the banning of Thalidomide – and not averted in Britain and Canada, where drug licensing processes were more expeditious – set a pattern for strict-and-slow drug regulation in the United States. Arguably, that approach, and the millions of quality adjusted life years (QALYs) that it sacrificed, only abated when the politically effective group of AIDS activists started to petition to make new drugs much more swiftly available.

No intermediate bets. Many high stakes decisions play themselves out over many years, whether it is successfully launching a high technology company or dealing effectively with the potential of global warming. Where there are severe policy disagreements, it might be worthwhile to require participants to make bets, at least reputational bets, whose resolution depends on intermediate markers. This is common in the private sector. If the high tech firm falls short of its business plan, it will have much greater difficulty raising the next round of financing. But in the policy arena, poor predictions usually get off completely free. One reason is that individuals are rarely required to make predictions/commitments about nonsalient intermediate outcomes, which though not striking may be highly informative. The greater is uncertainty, the easier is distortion, since substantial deviations can be dismissed as chance outcomes.

On a few occasions in the past, when parties were unable to agree on current actions in international environmental negotiations, it has been proposed to agree in advance on future actions to be taken contingent on specified future signals. The deals were never concluded, however, for two stated reasons. (1) They would not secure the promised actions without further bargaining and argument, because instead of fighting over what to do, parties would just fight over whether the triggering conditions had been met or not. (2) Pre-specifying the action-triggering contingencies too severely constrains the possibility of acting on the basis of real learning about the system – i.e., the signals we would choose now might well turn out to be the wrong ones. We believe, however, that it would be possible for parties to identify future contingencies as triggers of policy whose observation would be relatively (not entirely) uncontested, and that would not forestall future learning and adaptation – particularly if the promised actions are not extreme.

Some predictive implications. Given the bias of group processes towards the status quo, and its information-suppression accompaniments, we can predict a qualitative pattern when a group changes its mind some risk. Relevant examples are the dangers of cigarette smoking, the appropriate degree of pre-approval drug testing, or the actions to be taken vis-à-vis global warming. For long periods of time, information will accumulate that is contrary to established views and policies.

Partisan establishment of evidentiary thresholds for decisions

Common conditions of many such situations:

- evolving risk, the status quo is no response
- factions have different policy preferences and different prior views of the risk, but all observe the same evidence of risk, which varies with time;
- partisans of the status quo, i.e., opponents of action, largely get to control what level of evidence is required to act; they have this power because by the time the issue has become contentious at all, the activists believe present evidence already justifies action. The opponents, who say that it does not, are the ones normally obliged to say what further evidence they would regard as sufficient. Since activists are usually too committed to arguing that current evidence is more than sufficient, they do not engage in the debate over what increased level would be sufficient.
- opponents of action set the evidentiary threshold high, commonly invoking either the criminal-justice standard of proof, or the fabled "5% significance" of empirical research – i.e., the probability must be greater than 95% that the evidence indicates a real risk, before we act on it.
- the public and uninvolved legislators only respond when new stark information comes along, or a strongly organized pro-action constituency emerges.

Consequences

- Action is taken only when a very strong signal is observed, which is often too late.
- Moreover, in situations with slow dynamics, the action that is finally taken cannot deflect trends for some time. Consequently, when the evidentiary threshold is finally passed, it is likely that a further succession of increasingly alarming signals will quickly follow, independent of the action taken. With the SQ bias now overcome and the first action taken, a succession of increasingly extreme additional actions are likely to follow.
- We expect the following general pattern: the first response to an evolving risk will be taken too late – but once taken, is likely to be followed by more extreme evidence of risk and more extreme actions. We make no claim that the subsequent extreme actions will be sub-optimal, conditioned on the delayed start (i.e., the late start may require make extreme responses – it depends on the dynamics of the system and the form of damage function). We do claim that a preferable response for all would be to begin responding earlier with smaller actions, taken on less extreme evidence.

An Example: Ozone Depletion. The history of response to claims of ozone-layer depletion provides a canonical example of these dynamics and biases in action. Despite the gradual accumulation from the mid-1970s to mid-1980s of suggestive evidence that the threat was real,

opponents succeeded in resisting proposals to control the offending chemicals.⁵ Their argument that the decisive evidence that ozone depletion was a real risk would be detection of a statistically significant decline in global ozone came to be accepted by nearly all participants in the debate.⁶ But the ability to detect such a trend even if it was occurring was hampered by high natural variability of ozone in space and time -- including natural oscillations that were believed to occur over periods of more than a decade -- as well as by weaknesses of the observing network.⁷

Multiple analyses of ozone data conducted through 1985 found no significant trends (or in one case a small significant positive trend). A few analyses began reporting significant declines around 1985, but these were dismissed for various reasons – known problems with the instruments, a pattern of reported losses that did not match theoretical predictions, or losses reported only at certain altitudes. The shocking observation of the ozone hole, first reported in 1985, sharply raised concern but was not accepted as the required evidence of global depletion – principally because current stratospheric models could not account for it, and its cause was not known.

The required evidence of global depletion was only provided in 1988, by an international panel that re-analyzed all sources of ozone data and found statistically significant losses throughout the north temperate latitudes worldwide: 1.7 to 3 percent loss of ozone since 1969, larger in winter and at higher latitudes. This result galvanized the policy debates, provoking even long-time opponents to endorse phaseouts of ozone-depleting chemicals.

Over the next several years, reports of ozone losses grew rapidly, in the Antarctic, Arctic, and temperate latitudes.

After record Antarctic ozone losses were first publicly reported in 1985 (following record low values in 1979, 1981, and 1983 that were not previously reported), the Antarctic ozone hole repeatedly set new records for intensity and spatial extent – in 1987, 1989, 1990, 1991, 1992, 1993, 1995, 1996, 1998, and 2000.

Over the same period, progressively larger ozone losses were also repeatedly observed:

- in temperate latitudes (e.g., the estimated rate of loss was nearly doubled in 1991 from that reported in 1988; subsequent updates in 1992 and 1993 gave further increases in the estimated loss trend).
- and in high northern latitudes: (new records set in winters of 1990, 1991, 1993, 1995, 1996, 1997).

⁵ There were bans for their use as aerosol spray propellants in four countries. These limited controls, which were largely ineffective, were also nearly costless.

⁶ A small group of committed activists continued to argue that existing evidence – many laboratory and atmospheric measurements indicating the proposed loss mechanism was real, and model calculations projecting significant future losses – was already sufficient.

⁷ More generally, when there is a significant amount of background noise, more exaggeration can be expected in policy predictions, because error can always be dismissed as due to chance. Such a phenomenon is found, for example, in stock market analysts' predictions of corporate earnings. The greater is variability of earnings, the more such estimates are inflated. (Analysts are overwhelmingly on the sell side, hence have a bias.)

ABOUT THE NATIONAL SCIENCE FOUNDATION

The National Science Foundation (NSF) funds research and education in most fields of science and engineering. Awardees are wholly responsible for conducting their project activities and preparing the results for publication. Thus, the Foundation does not assume responsibility for such findings or their interpretation.

NSF welcomes proposals from all qualified scientists, engineers and educators. The Foundation strongly encourages women, minorities and persons with disabilities to compete fully in its programs. In accordance with Federal statutes, regulations and NSF policies, no person on grounds of race, color, age, sex, national origin or disability shall be excluded from participation in, be denied the benefits of, or be subjected to discrimination under any program or activity receiving financial assistance from NSF (unless otherwise specified in the eligibility requirements for a particular program).

Facilitation Awards for Scientists and Engineers with Disabilities (FASED) provide funding for special assistance or equipment to enable persons with disabilities (investigators and other staff, including student research assistants) to work on NSF-supported projects.

The National Science Foundation has Telephonic Device for the Deaf (TDD) and Federal Information Relay Service (FIRS) capabilities that enable individuals with hearing impairments to communicate with the Foundation about NSF programs, employment or general information. TDD may be accessed at (703) 292-5090, FIRS at 1-800-877-8339.

The National Science Foundation is committed to making all of the information we publish easy to understand. If you have a suggestion about how to improve the clarity of this document or other NSF-published materials, please contact us at plainlanguage@nsf.gov.