J Risk Uncertainty (2006) 33:131–149
DOI 10.1007/s11166-006-0175-8

Analyzing disaster risks and plans: An avian flu example

Baruch Fischhoff · Wändi Bruine de Bruin · Ümit Güvenç · Denise Caruso · Larry Brilliant

Abstract Narrative approaches to analyzing risks seek to identify the variables critical to creating and controlling a risk, then to instantiate them in terms of coherent themes (e.g., organizational failure, strategic surprise). Computational approaches to analyzing risks seek to identify the same critical variables, then to instantiate them in terms of their probability. Disaster risk analysis faces complex, novel processes that strain the capabilities of both approaches. We propose an approach that integrates elements of each, relying on what we call *structured scenarios* and *computable models*. It is illustrated by framing the analysis of plans for a possible avian flu pandemic.

Keywords Scenarios · Uncertainty · Risk analysis · Judgment · Avian flu

JEL Classification H89, I1, O21

Avian flu has appeared on the horizon like many other potentially disastrous threats, forcing already harried individuals to ask questions like: How big a deal is this? What do I need to know about it? Who are these self-proclaimed "experts"? How much do they know? What can I do about the risk, if officials do their job—and if they don't? Analogous questions arise with antibiotic-resistant bacteria, tsunamis, storm surges, financial collapses, improvised nuclear devices, weaponized anthrax (plague, tularemia, smallpox, etc.), cyber threats, invasive species, ebola, aging earthen dams, dramatic erosion of civil liberties, and environmental refugees, among possible disasters.

Answering these questions is objectively challenging. Each threat involves complex interactions among processes that are hard to understand in isolation. Available knowledge is scattered over scientific disciplines that rarely interact, much less assemble themselves into the teams needed to extract the knowledge that decision makers need. Experts may find themselves outside their comfort zones, given the novel circumstances, the need to interact with other disciplines, and policy makers' demands.

B. Fischhoff (✉) · W. Bruine de Bruin · Ü. Güvenç · D. Caruso · L. Brilliant
Carnegie Mellon University, Denise Caruso, Hybrid Vigor Institute, Larry Brilliant, Google Foundation & Seva Foundation
e-mail: baruch@cmu.edu

We offer an approach to analyzing such risks, whose execution and results should be within the reach of the organizations typically entrusted with managing disaster risks. It draws on what we will call *narrative* and *computational* analytical methodologies. It is informed by behavioral research regarding the individuals who perform analyses, who rely on their results, and who have their actions appear in the analyses (e.g., how they respond to evacuation instructions). It is illustrated with a case study of disaster planning for a possible avian flu pandemic.

1 Analytical philosophies

Approaches to understanding complex, uncertain processes can be arrayed along a continuum ranging from *computational* models to *narrative* scenarios. At one end, modelers attempt to identify relevant variables, characterize them in measurable terms, assess distributions of possible values, postulate dependencies between variables, and estimate their strength. Simulation runs sample a value for each variable, then predict the outcomes following from them. If these values are sampled according to their probability, the set of runs produces a probability distribution over possible outcomes. Planners who trust the model can use these results to focus their efforts. If they have access to the modelers, planners can also ask for new simulations, reflecting different assumptions about how the disaster will present itself or about how they will respond.

Each run of a computational model represents a possible future. However, simulations typically produce so many runs that none are examined in any detail. Narrative analyses sacrifice this breadth for depth, looking at a few possible futures in detail. The resulting scenarios are rarely quantified. Indeed, narrative analysts often explicitly caution against assigning probabilities to scenarios. Nonetheless, they often describe the work in terms that echo modeling (e.g., Schwartz, 1991). A scenario should address the "key factors" affecting future outcomes, connected according to their dependencies—tasks analogous to specifying the variables and relationships in a model. Each factor is represented by a specific value, akin to instantiating the variables in a model run. Those values are selected according to an organizing theme, sometimes a "major trend" (e.g., globalization) or a more modest "driving force" (e.g., population growth, increasing energy prices). These themes are designed to ensure that the values assigned to the factors could, in fact, co-occur, thereby passing a *coherence* test. In contrast, there is no guarantee that values chosen by their probability can co-occur. Indeed, best-case and worst-case scenarios are tools of computational analysis, rather than narrative analysis, unless the analysts have worked through how everything could conceivably go right or wrong. Scenarios create extreme cases by positing extreme themes (e.g., terrorists attacking a chemical plant under the cover of a Class 3 hurricane).

1.1 Narrative vs. computational analysis

Proponents of computational analysis claim that models can (a) accommodate diverse forms of knowledge in a transparent form, amenable to external review; (b) replace imperfect mental arithmetic with reproducible calculations; (c) allow explicit sensitivity analyses, revealing the implications of recognized uncertainties; and (d) focus research and action, by identifying the factors most worth understanding and shaping.

Proponents of narrative analysis claim that scenarios can (a) "stretch the mind," producing insights by the disciplined synthesis of otherwise scattered facts; (b) bound uncertainty, by showing the range of possible futures; (c) reduce complexity, by identifying interdependencies; (d) facilitate evaluation of potential actions, by placing them in a concrete setting; and (e) coalesce users, around a shared narrative (or set of alternative narratives).

Critics of computational analysis claim that models can (a) omit knowledge that is not readily quantified and forcefully represented in the analytical team; (b) be too complex to get the vigorous, independent review needed to reveal their flaws; (c) neglect the coherence of the futures implied by their model runs, thereby distorting the probabilities computed for possible outcomes; (d) limit their sensitivity analyses to the possibilities imagined by their analysts; and (e) fail to create the shared mental model that users need to coordinate their actions.

Critics of narrative analysis claim that scenarios can (a) frustrate external review, by obscuring which details are essential to their validity; (b) provide little guidance on when changes in the world (or in scientific understanding of the world) require updating (or abandoning) a scenario; (c) implicitly invite users to impute probabilities to them, even when cautioning against doing so; (d) invite exaggerating those probabilities, by the very act of giving scenarios narrative form; and (e) be so vague that users cannot be sure that they are all talking about the same thing or extract clear action implications.

In principle, the relative efficacy of narrative and computational analysis is an empirical question. In practice, though, a comparative evaluation is implausible. In circumstances important enough to merit serious analysis, no organization will allow itself to be randomly assigned to conduct narrative or computational analyses nor will it make equal investments in competing forms of analysis, then see which does better. That is true whether the organization is preparing for a disaster, looking for a strategic market opportunity, designing a political campaign, or developing any other high-stakes plan.

Rather than being driven by evidence, reliance on computational or narrative analysis seems to be determined by organizational processes. Models seem most common in organizations that prefer to have semi-autonomous experts develop solutions for them. Scenarios seem most common in organizations that prefer to have their own staff develop solutions through deliberative processes. The adoption of model-based solutions is aided by their clarity and the expertise of their creators. It is impeded by those solutions having evolved outside of organizational life, meaning that they lack an internal constituency that understands and embraces them. The adoption of scenario-based solutions is aided by the enthusiasm that group processes can generate. It is impeded by the vagueness of narrative accounts, meaning they fail to provide clear directions.

1.2 An integrated approach

Thus, computational and narrative analyses have complementary strengths and weaknesses. We propose an integrated approach that exploits this complementarity. It seeks to capture the social and expository value of scenario-based planning, while adding the rigor of model-based planning. In any application, its products will be a *computable model* and a set of compatible *structured scenarios*. A *computable model* captures the key variables and the relationships between them, with sufficient precision to be able to predict outcomes, were its data requirements satisfied. It should be clear enough that the relevant science can be mapped into it, as a step toward quantitative estimation. A s*tructured scenario* provides a narrative "run" of a computable model, instantiating each variable, with values chosen according to an organizing theme. It should be explicit enough to allow determining the compatibility (or coherence) of its values.

One could start the process with a model or a scenario. At times, interest in a risk arises from a vision of potential disaster; at times, it arises from a more abstract desire to know how a domain works. If one starts from a scenario, the next step is identifying its causal factors, dependencies, uncertainties, and outcomes. Those are then expressed in terms of a computable model, allowing review by subject matter experts. After needed revision, structured scenarios are created from the model, which are then evaluated for coherence, asking whether those futures could, in fact, occur. The process iterates until the modelers and scenarists are reasonably comfortable with their joint products. It could equally begin with a model created by computational analysts, with structured scenarios being used to assess how well its abstractions correspond to conceivable realities.

This process formalizes the checking that each approach aspires to do anyway. Narrative analysts naturally ask whether their scenarios capture the critical processes creating or controlling a risk; mapping scenarios into a computable model makes that test more explicit. Modelers naturally ask whether their runs represent plausible futures; instantiating some of them in concrete scenarios makes that test more explicit.

Once the process has been completed, either kind of analysis can follow its normal procedures. Thus, narrative analysts can create the text needed to bring structured scenarios to life, with greater confidence that they are faithful to the relevant science and that, once the scenario-focused discussions are done, clear plans can be generated. Computational analysts can estimate model parameters and run simulations, with greater confidence that they have captured the issues on planners' minds and that, once the calculations are done, their results can be conveyed to those who must act on them.

How far such analyses proceed should depend on their anticipated marginal utility. Any organization with disaster planning responsibility should be able to create a computable model, by requiring its experts to translate their knowledge into variables and relationships, then pool those beliefs into a common model. That effort may be all that an organization can afford, given the scarcity and expense of skilled computational modelers. A computable model and illustrative structured scenarios may also be all that an organization can absorb. As mentioned, scenarios' appeal comes from their accessibility to all members of an organization. The greater the quantification, the weaker the intuitive feeling that users will have over how the model works and the more they will have to take its results on faith. When decisions are made by senior policy makers, they may need quantitative estimates. With decision making distributed throughout an organization, its members may need a common model that they can share and update with experience. Allocating resources for disaster management should fall in the former category, benefiting from quantification. Coping with disasters once they arrive should fall in the latter category, benefiting from a widely shared narrative.

We have pursued the joint development of structured scenarios and computable (and sometimes computed) models in various contexts, including the risks and benefits of xenotransplantation (Güvenc, 2005), emergency contraception (Krishnamurti et al., 2006), sexually transmitted infections (Downs et al., 2004), dietary supplements (Eggers and Fischhoff, 2005), climate change (Casman et al., 2001), and radiological dispersion devices (Dombroski et al., in press). We will illustrate the process with a project in progress, coping with the potential disaster of an avian flu pandemic. These analyses have been produced at low cost, with mostly donated labor from various subject matter experts. Although these models and scenarios set the stage for more detailed computational and narrative analyses, we believe that even this initial stage provides a useful picture of this potential disaster.

2 Avian influenza

2.1 Disaster planning context

In early October 2005, 49 individuals were invited to Pandefense 1.0, a closed meeting scheduled in a month's time, with the goal of evaluating non-pharmacological interventions to reduce the threat of a possible avian flu pandemic. The terms of the invitation preclude revealing the names of participants, other than the organizers. All were asked to complete a survey that would provide a snapshot of their views, to be presented at the meeting (Bruine de Bruin et al., 2006). Of the 36 who completed the survey, 19 described themselves as medical experts, with specialties mostly related to epidemiology and influenza. The 17 non-medical experts represented sociology or political science (6), business (6), disaster-relief non-profits (4), and politics (1).

A set of scenarios central to US pandemic planning was used to create a computable model. That model structured the survey, so that responses to its structured questions would provide preliminary estimates of model parameters, while responses to its open-ended questions would suggest refinements to the model's structure. The model was revised in the light of survey responses, peer review at Pandefense 1.0, and other expert input. At various stages, it was translated into structured scenarios, as a check on its completeness and coherence.

Table 1 presents an excerpt from one of several pandemic disaster scenarios created by the Centers for Disease Control (Harris, 2005). Each of these scenarios addressed all variables that its creators thought critical to understanding a pandemic's evolution and impacts, as well as the effects of possible disaster plans. The value assigned to each variable is meant to be consistent with the theme embodied in each scenario's initial conditions. The dependencies among the variables reflect the analytical team's understanding of the relevant science (e.g., the death rate among those infected, the social conditions creating gray markets). The other CDC scenarios embodied the themes following from other initial conditions.

Such a narrative analysis seeks to ensure coherence by making the events consistent with one another and the initial conditions. However, no probabilities are assigned to either the initial conditions or the contingencies. As a result, it is unclear how seriously to take any of its details. For example, panic is very unusual in disasters (Wessely, 2005), despite its popularity in the popular mind and among non-social scientists (Fox, 2006; Levi and Kelly, 2002). The scenario's prediction of social disorder could reflect a fundamental misunderstanding that undermines its overall credibility. Or, it could reflect a deliberate choice to interject a dramatic element, which is not to be taken seriously. A computable model makes such assumptions explicit—and open to peer review.

In order to create such a model, we identified the variables and relationships that appeared to be the driving forces in the full version of the scenario in Table 1. These were assembled into a draft model, into which two other CDC scenarios were mapped. When the fit was poor, the model was revised, so that all factors that CDC believed to be relevant were accommodated. As a tribute to their creators, these scenarios addressed similar variables, even when one could not clearly ascertain the relationships among them. The model sought to express the issues in CDC's scenarios in terms that corresponded to the relevant research literatures. The revised model was then supplemented with other factors drawn from the research literature (some outside CDC's scope), as well as input from the survey, Pandefense 1.0, and other experts.

Table 1 Excerpts from an (unstructured) avian flu scenario

In late June, the CDC reports that the virus has been isolated from ill airline passengers arriving in four major United States cities. . . Vaccine manufacturers are requested to shift vaccine production from annual to pandemic vaccine.
In July, small focal outbreaks begin to be reported throughout the United States. The first doses of a new pandemic vaccine become available in September. Despite full-scale production by manufacturers, supply remains very limited. . . Community-wide outbreaks begin to occur more frequently as children return to school, and by late August, outbreaks are occurring simultaneously throughout the country. . .
Overall, about 2 percent of Americans with influenza illness die. In communities during the peak weeks of. . . outbreaks, about a quarter of workers are absent because of illness, the need to care for ill relatives and fear of becoming infected.
Hospitals are overwhelmed and staff shortages limit capacity. Intensive care units at local hospitals are unable to provide care for all who need it, and there are shortages of mechanical ventilators for treatment of patients with severe pneumonia. Makeshift hospitals established in schools and armories care for those who are unable to be treated in regular hospitals. . .
During the peak of disease activity in the community, police, fire and transportation services are limited by personnel shortages, and absenteeism at utility companies leads to spot power outages. Supplies of food, fuel and medical supplies are disrupted as truck drivers become ill or stay home from work.
In some areas, grocery store shelves are empty and social unrest occurs. Long lines form where food and gasoline are available. Elderly patients with chronic, unstable medical conditions hesitate to leave their homes for fear of becoming seriously ill with influenza.
Riots occur at some vaccination clinics as people are turned away or supplies run out. Several trucks transporting vaccine are hijacked, and a gray market develops for vaccine and antiviral drugs—many of which are counterfeit.
Pig herds acquire infection with the pandemic virus and are decimated; large numbers of workers in those settings become ill.
Family members are distraught and outraged when loved ones die within a matter of a few days. Public anxiety heightens mistrust of government, diminishing compliance with public health advisories. "Worried well" seek medical care despite their absence of influenza illness, further burdening the health care system.
Mortuaries and funeral homes are overwhelmed.
(Source: Harris, 2005)

2.2 Modeling pharmacological responses

Figure 1 shows part of the model, addressing two pharmacological interventions, whose availability will determine the role of non-pharmacological ones: *vaccines* and *anti-virals*. The modeling language is based on that of influence diagrams (Clemen, 1997). Ovals are uncertain variables, which need to be predicted. Rectangles are actions, which need to be planned and implemented. Arrows are "influences," in the sense that the value of the variable at an arrow's tail should influence predictions of the variable at its head. For example, knowing the details of a *vaccine strategy* should influence predictions of the associated *healthcare costs*. Those details, together with estimates of a vaccine strategy's *efficacy* and the disease's *rate of spread*, should influence predictions of *morbidity* (from the flu itself and associated illnesses). The five gray ovals are the focal outcomes of the analysis: *morbidity, mortality, healthcare costs, non-healthcare economic costs*, and *social costs*.

The next step toward computability is specifying the outcome variables, first in terms that capture their underlying rationale, then in measurable form. That specification should improve communication between model producers and model consumers, while identifying needed data and expertise. For example, *social costs* might be defined as "disruptions of everyday life that are not readily monetized, including pain and suffering, dislocation

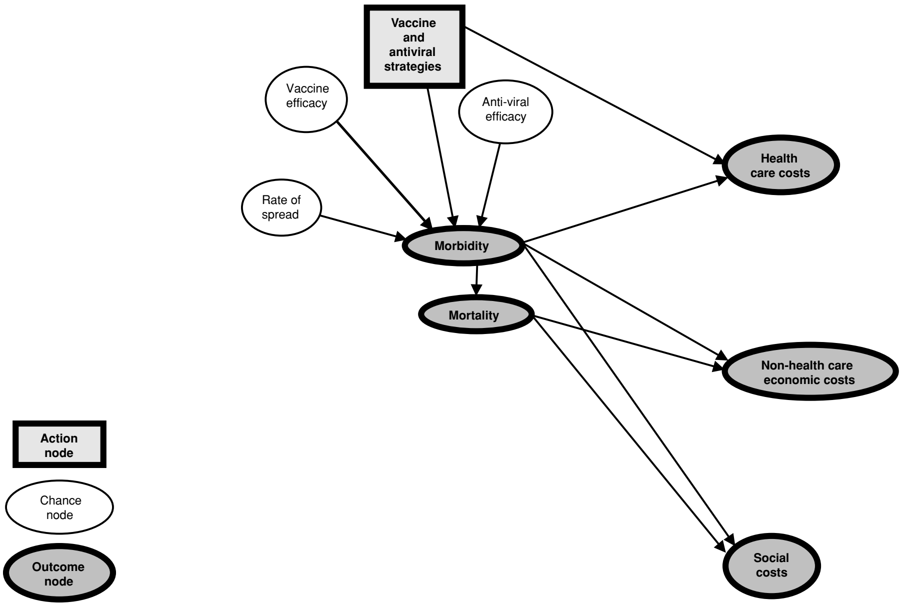

Fig. 1 Top-level risk model for pharmacological interventions and potential effects. Ovals indicate uncertain variables, which need to be predicted. Rectangles indicate actions, which need to be planned and implemented

of families and communities, stress and potentially post-traumatic stress disorders, loss of access to education due to school closures, fear and distrust due to misinformation, vigilantism, etc." That specification may translate a "soft" variable into terms that afford it status more equal to that of "hard" variables (e.g., *morbidity*, *mortality*, *healthcare costs*). There are technical conventions for defining some of these variables (e.g., *rate of spread*), while others have varying definitions (e.g., *healthcare costs* might, or might not, include nursing homes, home health, and insurance). As discussed below (under "risk characterization"), these specifications inevitably embody social values. One benefit of computational modeling is bringing those issues into relief.

Creating a computable model requires no more (and no less) than clear thinking about the precise issue that each node and link is meant to express. Table 2 shows an auditing procedure for nudging models toward computability, which we have used in helping non-scientists to create computable models. Although second nature to scientists, these questions are sufficiently straightforward that non-scientists can use them to check their thinking. If so, then any group charged with disaster planning could create a model organizing its knowledge about the processes creating and controlling those risks. Those members who intuitively think about risks in narrative terms should be comfortable with the verbal formalisms of a computable model, while those who prefer computational analyses should feel that such rigor has not been abandoned, even if the organization's discourse is in terms of structured scenarios.

2.3 Behaviorally realistic analysis

Even this small model shows the three forms of behavioral realism needed for any risk management, including disaster planning:

Table 2 Questions to ask in checking a model's clarity

Node review

Complete the following for each node:
1. Name of variable (or vector of related variables)
2. Possible values of the variable(s)
3. Possible procedures for measuring variable
4. Methods for measuring variables

Single link review

Complete the following for each link:
1. Names of nodes involved.
2. Simple statement of the link (e.g., X causes Y because; X is a good indicator of Y because).
3. If there are multiple variables at a node, does this simple statement hold for each combination of variables? (If not, consider partitioning the variables into separate nodes.)
4 Source and strength of claim for link. (Use dashed lines for speculative links or ones whose existence is in dispute.)
5. (optional) Strategies for studying link.
6. (optional) Strategies for affecting link.

Multiple link review

Complete for each link:
1. Does it go into a node that also has only one link going out? (If so, the intermediate node could be eliminated, unless having it provides a useful reminder of the connection between the nodes that it separates.)
2. Does it have the same input and output arrows as another link? (If so, consider combining them or representing that area in the influence diagram as a single topic in a higher-order [simpler] model.)
3. Is it part of a circular chain of links? (If so, identify the time dependency among the links—or group the chain in a single node, with its own internal dynamics.)

Overall model review

1. Are critical endpoints easily identifiable?
2. Would connecting any pair of unconnected nodes add predictive value?
3. Is there feedback from the endpoints to the initial conditions (indicating temporal dynamics)?
4. Are there important "index variables" that affect many model values, within the basic structure (e.g., gender: for a disease with different expressions for men and women)?

2.3.1 Recognizing the roles of expert judgment

A model shows the expertise that planning requires. Some of its relationships may have been studied extensively enough to allow quantitative predictions. For example, there are standard procedures for estimating the *healthcare costs* of a *vaccine strategy*, once it has been specified in terms of how the vaccine will be developed, manufactured, stored, distributed, and administered (Kaufmann et al., 1997). The research base for other variables allows less precise predictions (e.g., the *social costs* of widespread *mortality*). Morgan and Henrion (1990) offer a protocol for eliciting expert judgments that is informed by studies of the cognitive processes involved. Morgan and Keith (1995) apply it to climate change.

Using expert judgments effectively requires understanding their *internal validity*, in the sense of the soundness of the relevant research, and its *external validity*, in sense of how well it generalizes to the specific setting. Funtowicz and Ravetz (1990) propose characterizing internal validity in terms of four features: (a) *proxy* (how close a field's observable measures are to actual phenomena), (b) *empirical* (how solid its evidence is); (c) *method* (how well established its procedures are), and (d) *validation* (how much independent confirmation it has). Krayer von Krauss et al. (2004) provide an application to GM crops. Once a computable model has been created, its nodes and links can be characterized in these terms, showing users the quality of the science potentially available. That knowledge allows them to understand what uncertainty they face and how to allocate resources to its reduction.

A computable model can also structure assessing the external validity of available research. Looking "upstream" from a model element, one can ask whether studies of that issue have typically involved conditions differing from those of the application. Looking "downstream" from a model element, one can ask how well the research has addressed those contingencies. Disaster planners can err by assuming that existing research can be applied directly to novel circumstances—and by assuming that "everything is so different" that they can ignore the research and substitute their own intuitions.

2.3.2 Recognizing the roles of human behavior in the analysis

The evolution of any disaster depends on people's behavior (e.g., will they take a novel vaccine, will they self-medicate or go to emergency wards, will they work while ill). As a result, a model should reflect the relevant behavioral science, as much as the relevant natural science. In practice, though, that is rare, unless the behavioral science has a computational form, like the accounting formulae for *healthcare costs*.[1] However, even such formulae may capture only a portion of the relevant behavior. The current problems with the Medicare Part D prescription benefit plan reflect limits to purely economic analysis. It would be similarly naïve to use reported case-fatality (*mortality/morbidity*) rates, without knowing how well *surveillance* processes identify and report the dead and the sick, including those who are asymptomatic or hide their symptoms in order to avoid social sanctions.

Behavioral variables are largely absent from Figure 1. As such, the model reflects current US policy, whose investments are mostly in drugs (Salaam-Blyther and Chanlett-Avery, 2006; Department of Health and Human Services, 2006). When those programs are implemented, they will have concrete form, each of whose elements (development, manufacture, storage, etc.), will depend on behavior, so that the plan itself could be treated as an uncertain event (i.e., changing the rectangles to ovals) (Brown, 2005).

Narrative analysis is less vulnerable to the self-censoring common to computational analysis, whose practitioners must always worry, "where will I get the data?" As a result, basing computable models on structured scenarios reduces the risk of excluding behavioral variables that are not readily (or customarily) quantified. Even if those variables are omitted in subsequent computations, that will be done with an awareness of the gaps. Figures 2 and 3 (below) show the analytical domain implied by the CDC scenarios and the deliberations following from them.

2.3.3 Recognizing analysis users' informational needs

Figure 1 focuses on a variable (*healthcare costs*) that is vitally important to the professionals and politicians who must raise and allocate the relevant funds. That variable might be less important to citizens who expect their society to spend whatever it takes to protect them against a rare, grave threat. For them, analyses based on Figure 1 might ignore such critical issues as the effectiveness of self-protection measures, the robustness of community services (e.g.,

[1] In June 2006, the NIH's Models of Infectious Disease Agent Study (MIDAS) program and the Brookings Institution convened a workshop to address just this problem, in the context of pandemic planning.

police, sanitation), and responses to absenteeism—all issues raised by CDC's scenarios. In order to address these needs, analysts need to identify the decisions faced by individuals, in the context of the impending or exploding disaster, as defined by the model and scenarios. Doing so for every individual and decision would quickly exhaust analytical resources. However, it may be feasible to analyze, in general form, the decisions faced by classes of individuals (e.g., parents of small children, the immunocompromised, the uninsured) (Fischhoff, 2005). With the fateful choices posed by disasters, the return on that investment may be large, both in absolute terms (considering the ad hoc advice that it would replace) and relative to fine-tuning analyses for policy makers.

The usefulness of any analysis depends on how well its results are understood. A computational model provides the structure needed for effective communication. It defines terms precisely enough that analysts can determine whether they have been understood as intended. For example, individuals concerned about the *efficacy* of pharmaceuticals might not know how experts would define that term. That confusion might involve the intended dose (currently, much higher for avian flu vaccines than others) or the kind of protection provided (currently, quite limited for *anti-virals*). A model also identifies the context within which individuals interpret individual facts. For example, people might erroneously believe that there are *healthcare costs* associated with *mortality*, a link that is missing in the figure because corpses rarely cause illness (WHO, 2005). Such confusions are typically easy enough to correct, if one has studied the inferences that communications evoke. Morgan et al. (2001) summarize research using computable models to design and evaluate risk communications, sometimes using structured scenarios to enhance recipients' mental models (e.g., Downs et al., 2004).

Such communication research can also improve analyses' behavioral realism by providing estimates for the degree of compliance with policies (see Figure 3 below). For example, in the midst of the abortive 2002–2003 smallpox vaccination campaign, most Americans did not know that the vaccine was effective if received after exposure, but before symptoms appeared. A year after the anthrax episode, most people did not know that it was not contagious (Fischhoff et al., 2003). Plans that assumed otherwise would have been mistaken.

2.4 Scenario-based refinements

In a model, both the presence and the absence of links are informative. For example, Figure 1 predicts *non-healthcare economic costs* and *social costs* separately from *morbidity* and *mortality*—indicating that *morbidity* influences both outcomes, independent of its effects on *mortality*. Thus, knowing the case-fatality rate does not extract all of the predictive value of *morbidity*. In data-rich environments, one could use structural equation analysis to establish the need for these links (Burns and Clemen, 1993). In data-poor environments, one needs to think through the connections. Scenarios can structure that process, by making contingencies more concrete.

Consider, for example, the cost and benefit data that healthcare economists have accumulated from past pharmaceutical interventions. Those estimates should include the effects of predictable surprises, such as 2004's loss of a major flu vaccine production facility and 2005's chaotic vaccine distribution. However, when applying data to a new setting, one must ask whether the historical relationships hold. An avian flu scenario could "just" reflect extreme values from the historic distributions (e.g., the highest *morbidity* ever seen) for which some, sparse data exist. However, the scenario could also depict an unprecedented disaster, in which all relationships must be reconsidered. For example, a pandemic could disrupt supply chains so badly that vaccine distribution is severely hampered. Table 1's scenario raises such prospects in general terms; a structured version would frame them more precisely. This scenario does not consider the discontinuities possible with changes like legislation that reduces vaccine manufacturers' liability, new vaccine technology, or the waiving of testing requirement in a pandemic. Narrative analysis could create scenarios addressing these possibilities, leading to appropriate refinement of the model.

There are three ways to address such possibilities, as a model moves towards computation. One is by adding uncertainty, for example, flattening the *healthcare costs* distribution to accommodate the possibility of much higher costs (e.g., from distribution breaking down) or much lower costs (e.g., from a technological breakthrough whose savings are shared). The second is by adding nodes, such as one for *vaccine availability* between *vaccine strategy* and *healthcare costs*. These nodes would have their own predictors and experts (e.g., labor relations specialists familiar with issues like those that undermined the 2003 smallpox vaccination campaign). The third is treating the vaccine program as an uncertain variable, recognizing that it may emerge very differently than the plan's description of who will do what when (Brown, 2005).

Analogous questions can be asked about how each connection would play out, should a scenario come to pass. Just as the scenario tests the model, the model tests the scenario, by seeing whether it fits the model at all or implies implausible connections.

2.5 Fuller models

Although pharmacological interventions are a natural response to health problems, neither *vaccines* nor *anti-virals* are currently feasible strategies. Nor did either seem feasible in the next three years, according to Pandefense survey respondents (see below; Bruine de Bruin et al., 2006). Without unprecedented international mobilization, these strategies might never be

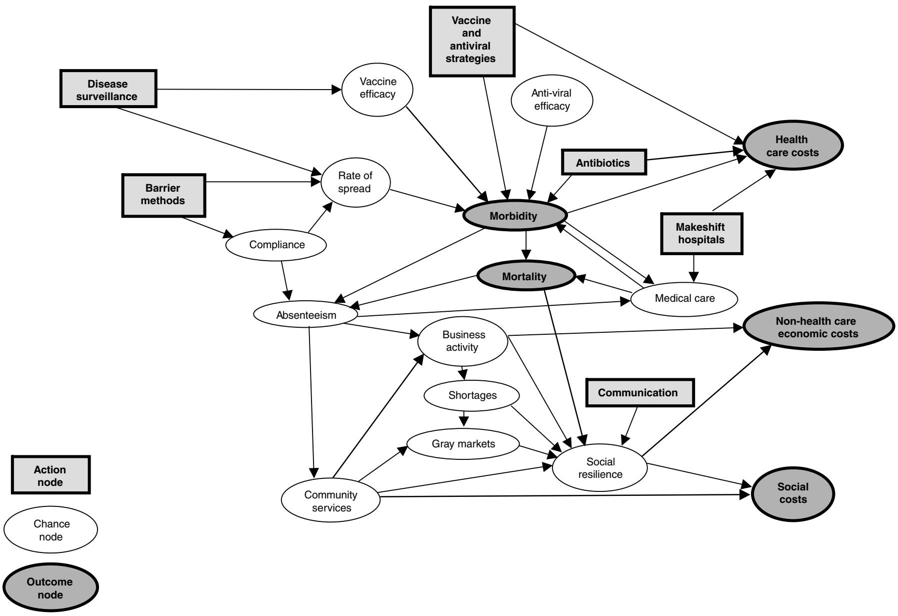

Fig. 2 Top-level risk model for a behavioral intervention, the use of barrier methods, along with intermediate processes affecting its impacts on the focal consequences of Figure 1. Ovals indicate uncertain variables, which need to be predicted. Rectangles indicate actions, which need to be planned and implemented

feasible for poor countries. Figure 2 expands Figure 1 to include key non-pharmacological interventions and some of the uncertain factors shaping their effects. These measures are (a) *barrier methods*, such as gloves and face masks, designed to keep society running by allowing healthy individuals to go about their normal lives (Institute of Medicine, 2006); (b) *antibiotics*, designed to reduce *morbidity* from secondary infections; (c) *makeshift hospitals*, bringing *medical care* to schools, fire stations, etc.; (d) *communication*, designed to enhance *social resilience*; (e) *disease surveillance*, tracking new cases.

Like other computable models, Figure 2's potential value is in giving analytical expression to issues raised verbally in CDC's scenarios, responses to the Pandefense survey's open-ended questions, and other sources. It pools those issues in a shared model, which sets the context for both narrative and computational analyses, informed by the science relevant to its variables and relationships.

Figure 2 shows *barrier methods* as affecting the *rate of spread*. Those methods' success depends on individuals' *compliance* and on *disease surveillance*, which determines how much warning citizens and officials have for implementing plans (and for formulating, producing, and distributing vaccine, thereby determining *vaccine efficacy*). Reading from the left, the figure shows the cascade of effects following from the interventions. The better the *compliance* with *barrier methods*, the lower the *absenteeism* rate, because people can go to work with an acceptable level of risk. The less *absenteeism* there is, the greater will be (private sector) *business activity* and (public sector) *community services* (e.g., police, fire, mail, sanitation). Reduced *business activity* would mean *shortages*, perhaps increasing *gray markets* (preferential treatment for the well-connected). Better *community services* would reduce *gray markets*. Many factors will affect the *social resilience* needed to ride out a pandemic, as well as the associated *non-healthcare economic costs* and other *social costs* (e.g., reduced public morale, faith in government). These processes also depend on the disease's progress, through its effects on *morbidity* and *mortality* and their influences on *absenteeism*, *medical care,* and *social resilience*. The effects of *antibiotics* on *morbidity* will depend on whether the flu kills people so quickly that they have no time to get secondary infections. Morbidity and mortality will depend on the quality of *medical care*.

Table 3 shows a structured scenario based on Figure 2's computable model, using values derived from the research literature and the Pandefense survey results. Its general form resembles that of the CDC scenario (Table 1), which was our point of departure. However, its specifics reflect the iterative analysis described above.

Figure 2 analyzes pandemic risk from a societal perspective, anticipating the conditions that any decision maker will face. How fast will people be dying? How sick will the survivors be? How much lead time will surveillance provide? Is the whole problem big enough to merit planning? Is enough known to plan with any confidence? In order to make specific plans, decision makers will want to know a lot about specific details. Some will want to know the factors affecting *absenteeism* and the opportunities to reduce its effects (e.g., telecommuting, enterprise health programs). Others will want to know the factors affecting *shortages* and the effects of supply-chain adjustments.

Figure 3 expands on the processes influencing a variable critical to the success of any behavioral intervention, *compliance*. Unlike Figures 1 and Figure 2, which are derived primarily from pandemic-related public health sources, Figure 3 is based on basic behavioral science (e.g., Fischhoff et al., 1998), which public health documents treat sketchily. It posits three primary predictors for *compliance*: (a) how well people understand the strategy (*comprehension*), (b) how credible it seems as a way to reduce the risk (*trust*), and (c) how well it can be executed (*feasibility*). Figure 3 shows *comprehension* and *feasibility* as influencing *trust*, anticipating that people are more likely to trust a program that they understand and

Table 3 A model-based scenario

(Figure 2 model variables italicized, in brackets when synonyms used in text)
Imagine that H5N1, the virus that causes avian flu, has become transmissible from human to human [*rate of spread*]. *Vaccines* and *anti-virals* are not available in sufficient quantities to stop the pandemic. Within a few months, as many as 100,000,000 in the United States may get sick with this human form of bird flu [*morbidity*]. As many as 6,000,000 may die [*mortality*].
Many sick people seek *medical care*, dramatically increasing *healthcare costs*, and so overwhelming hospitals and clinics that many people do not receive needed *medical care*. As a partial stopgap, *makeshift hospitals* are created in places like schools and fire halls.
Patients with avian flu often get secondary infections. Some, like bacterial pneumonia, can be cured with *antibiotics.* However, buying and distributing them has gotten expensive [*healthcare costs*].
The outbreak also has massive *non-health economic costs*. Sick people [*morbidity*] stay home from their jobs [*absenteeism*]. So do healthy people, afraid of getting sick [*absenteeism*]. As a result, *business activity* is greatly reduced, creating *shortages* of food, fuel, and other goods. *Community services*, like gas, electricity, and sanitation, are also less reliable.
These *shortages* create *gray markets*, in which goods and services are available mainly to wealthy and well-connected individuals. These disruptions undermine *social resilience*, already challenged by death and illness [*mortality* and *morbidity*].
Government officials rigorously follow a *communication* strategy of describing the situation honestly, however grim those reports are. The feeling that everyone is "in it together" boosts *social resilience* and limits the *social* and *[non-health care] economic costs* of the strains on people's lives.
Barrier methods, like wearing masks, can reduce the disease's *rate of spread,* perhaps enough to reduce *absenteeism*, and improve *business activity, community services*, and *medical care.* Unfortunately, initial *communications* are so confused that some people get sick, either because they use the wrong kinds of masks or use the right masks in the wrong way. As a result, *compliance* with requests to use masks is low.

can execute. Each variable has its predictors. Some show additional influences of factors in Figure 2. For example, *gray markets* will affect *trust* in any official communication, as well as the availability of *supplies* that determine the *feasibility* of *compliance*. *Makeshift hospitals* and *antibiotics* appear as influences on *medical care* and *morbidity*, because of their effects on *feasibility*, which might be harder for someone who is ill (*morbidity*) or has *care-giving needs*. *Social resilience* appears with its Figure 2 influences and with a direct link to *trust*, implying predictive ability beyond its indirect links.

Figure 3 also adds interventions, suggested by these influences on *compliance*. *Household subsidies* would provide material resources to those whose *household finances* preclude acquiring the *private supplies* needed for *feasibility*. Such a program should indirectly increase *trust* by enhancing *social resilience*, through its demonstration of commitment to the poor. As elsewhere, predictions depend on program specification (who gets what, when). The model changes *community services* from a variable, affected by other processes, to a program, designed to enhance *business activity* and *social resilience*. If successful, its indirect effects could include enhancing *feasibility* and *trust*, vis-à-vis a *behavioral strategy*. Finally, the model identifies three kinds of informational programs: (a) *official communication* about the strategy; (b) *prior communication*, designed to shape *unofficial information*; and (c) *prior education* designed to shape the *prior beliefs* that determine *comprehension* of *official communications*, in the light of *unofficial information*.

Psychologists and others have studied *compliance* in many domains, providing a basis for predicting the processes on the left side of Figure 3. Sociological, economic, and other research provides a basis for predicting the processes on the right. No responsible plan for managing these potentially disastrous risks should ignore that research. It is hard to conceive

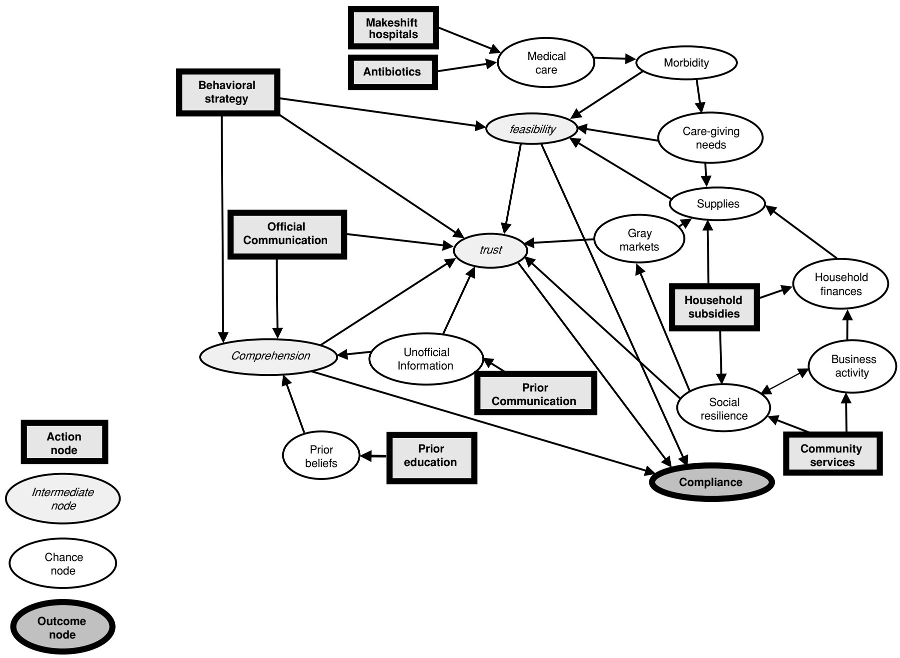

Fig. 3 Second-level model for factors shaping compliance with a behavioral intervention. Ovals indicate uncertain variables, which need to be predicted. Rectangles indicate actions, which need to be planned and implemented

of a scenario in which that research would be irrelevant—and even harder to imagine that useful speculations could come in its stead.

3 Using a computable scenario-based model

3.1 Knowledge management

Having a computable model allows systematically recruiting relevant evidence, ensuring that each node and link receives some attention, before examining any of them in depth. It also provides a template for organizing the flood of potentially relevant information (Fischhoff, 2000). On the day that this paragraph was first drafted, CDC's Immunization Issues in the News had 85 items on avian flu, the public health blog "Effect Measure" (http://scienceblogs.com/effectmeasure/) had 8 long comments on estimating the case-fatality rate's denominator (given the limits to surveillance), and Google retrieved 93 million hits on "bird flu" and "avian flu."

The standard for aggregating quantitative estimates is a rolling meta-analysis, for evidence of a single kind, or the Bayesian equivalent, for disparate types of evidence. However, even without formal aggregation, just collecting studies in decision-relevant form should reduce the risk of major oversight. As mentioned, anyone considering *social resilience* needs ready access to research showing how rare panic is (Wessely, 2005). Anyone considering an *anti-viral strategy* might benefit from Zamiska's (2006) discussion of the risk of prior use reducing those drugs' effectiveness in a pandemic. A computable model both flags those issues and provides a way to organize the evidence.

3.2 Risk characterization

Any analysis expresses a political-ethical position, most obviously in its choice of outcomes to predict. The gray ovals represent the effects determined to matter here, with intermediate variables chosen for their value in predicting those outcomes. *Compliance* (the focus of Figure 3) is a predictor for some people, an outcome for others (e.g., health officials evaluated by how well they implement a barrier program). Values are further expressed in how variables are measured, sometimes called "risk characterization" (National Research Council, 1996). For example, mortality measures must either consider or ignore the ages of the dead (Crouch and Wilson, 1981; Fischhoff et al., 1981), perhaps treating all deaths equally, perhaps assigning extra weight to deaths of the young, the old or the middle aged (Viscusi and Aldy, 2003). Quality Adjusted Life Year (QALY) measures take the ethical position that one should also consider the nature of the lives lost (or saved). Figures 1 and Figure 2 highlight *social costs* as an outcome, lest *non-healthcare costs* neglect ("intangible") effects that are not readily monetized (e.g., pain, suffering, stress without overt physiological damage). If martial law were considered as a *behavioral strategy*, then the ethical question would arise of whether to consider civil liberties in the definition of *social costs*.

3.3 Strategy design and evaluation

Expressing interventions in model terms allows specifying the processes by which they are expected to achieve their intended effects, as well as other processes affecting those outcomes and any possible side effects. For example, according to Figure 1, the impacts of investing in an *anti-viral* strategy depend on its *effectiveness* and the disease's *rate of spread*. A plan that fails to address these issues well warrants little trust. Figure 3 offers a template for analyzing the social factors that might support or undermine a behavioral strategy, perhaps suggesting ways to improve it or perhaps reasons to abandon it. For example, the US federal government's current disaster planning features "expectations management," in the sense of promising little direct help, while encouraging self-reliance. The model shows what it would mean to translate that theme into a well-structured plan. For example, if it precluded *household subsidies*, what assumptions must be made about *household finances* and *gray markets*, in order to ensure the *supplies* that people need to achieve *feasibility* for any *behavioral strategy*?

3.4 Communication design and evaluation

Evaluating plans requires predicting their effects on the behaviors in the model. As depicted in Figures 2 and Figure 3, *compliance* with behavioral interventions depends on people's *comprehension* of the instructions and *trust* in those dispensing them, as well as the *feasibility* of their demands, given the context within which they are advocated. As mentioned, the model also provides a framework for conveying the kind of integrating mental model that people need to understand a complex situation, making sense of competing claims and adapting their plans to changing conditions (Morgan et al., 2001).

4 Applying scenario-informed computational models

The survey reported by Bruine de Bruin et al. (2006) takes some next steps toward developing behaviorally realistic policies, based on this foundational work. Specifically, it elicits expert judgments precise enough to be used in computational models for one critical scenario: that in which a pandemic occurs within the next 3 years. Several examples will illustrate how formal policy analyses could build on these results.

4.1 How imminent is the threat?

As described in the structured scenario and model, a pandemic's evolution depends on the effects of interventions like *vaccines, barrier methods*, and *makeshift hospitals*. Those impacts depend, in turn, on both the interventions' inherent capabilities and the time available to implement them. In order to estimate that time, the survey's first question asked, "What is the probability that H5N1 or a similar virus will become an efficient human-to-human transmitter (capable of being propagated through at least two epidemiological generations of humans) sometime during the next 3 years?" The response options were: 0%, <1%, 10%, 20% , . . . , 100%.

Among the medical experts, the median judgment was 15% (range = [<1%, 80%]. The non-medical experts, recruited because they would play other roles in pandemic planning (e.g., logistics, telecommuting), saw a significantly higher probability (median = 60%; range = [30, 70%]). As a way of assessing the validity of these judgments, as capturing respondents' beliefs, the survey's final question asked how many years would have pass before transmission had a 10%, 50%, 90%, and 100% chance of occurring. These judgments were generally consistent with those for the three-year probability. The medical experts said that it would take 3 years until there was a 10% chance; the non-medical experts said that it would take 3 years until there was a 50% chance.

Thus, the two groups of experts would interpret pandemic scenarios with very different expectations of their imminence. Assuming that the medical experts are better informed about these issues, the different beliefs in the two groups suggest a communication failure, of a type that is common when no one produces and disseminates clear, quantitative summaries of expert opinion (Morgan et al., 2001). The probability of transmission affects pandemic risk in two ways: the chances of it happening at all and how it evolves. Analyzing that evolution requires examining a scenario's implications for each model variable. The next section considers some of those dependencies.

4.2 How feasible are the strategies?

The medical experts saw almost no chance of effective vaccines or antivirals being available in the next three years (median = <1%, for each). The non-medical experts saw significantly better chances (median = 15%, 30%, respectively). Thus, here, too, medical experts' beliefs had not been communicated to non-medical experts. Based on the medical experts' beliefs, no analysis should assume the availability of effective pharmacological interventions, in the next three years. That assumption should affect the probability afforded to any scenario and the value given to any model variable.

As mentioned, the US is currently investing almost exclusively in vaccines and antivirals. To the best of our knowledge, that strategy has not been systematically evaluated, nor compared with possible alternatives. Such analysis should begin by creating a structured scenario for what would happen in the (unlikely) event that a pandemic arrived in the next three years—and the current plan is followed. The model corresponding to that scenario would have very low probabilities for the availability of vaccines and antivirals, as well as low estimates for the effectiveness of most other interventions—given the lack of investment in them. For example, the analysis should assume greater *morbidity* and *mortality* than would be experienced with a national strategy of stockpiling *antibiotics*, for treating secondary infections of pandemic flu, or preparing *makeshift hospitals*, for supplementing regular hospitals that cannot provide adequate *medical care* for pandemic flu or routine illnesses (e.g., diabetes). *Social resilience* might be undermined, if citizens felt that government had failed to make seemingly obvious preparations.

The structured scenario's implications for each model variable should be similarly analyzed in narrative terms—then documented in a way that allowed others to examine the analysts' assumptions. Each variable could then be estimated quantitatively and the computational model run, to produce outcome estimates. Combining those estimates with the 15% chance of a pandemic gives the current national strategy's expected risk, in the next three years. The same analytical process would be applied to other possible strategies. Following it might reveal that one needs to go no further than creating structured scenarios and extracting their implications for the computable model—if those analyses reveal dominating or dominated strategies. If not, then further quantification is needed, proceeding until policy makers receive a clear enough picture to choose among strategies. The approach proposed here should allow that computation to proceed, without losing the trust of individuals who are inclined to narrative analysis—or their ability to contribute to the design of better strategies. Those analyses might find that our best gamble is to emphasize vaccines and antivirals, given the expectations for the pharmaceuticals working and the pandemic tarrying. However, those analyses do not seem to have been done. The stakes riding on them are seen in survey respondents' median best-case and worst-case estimates of 15 million and 100 million Americans sick, and 0.5 million and 6 million dead—in the event of an outbreak in the next three years, without sufficient vaccines and antivirals.

5 Conclusion

The combination of narrative and computational analysis proposed here accepts the reality of people's natural ways of thinking. By iterating between computable models and structured scenarios, it seeks to enrich the thinking of people who are most comfortable with either mode of analysis. That triangulation process should reduce the overconfidence that either models or scenarios can create in isolation. Creating structured scenarios should diminish any unwarranted advantage for model-friendly evidence. Creating computable models should diminish any unwarranted advantage for rhetorical claims. The combination is also sensitive to the role of communication in disaster planning and helpful for designing disaster communications.

A computable model provides a platform for integrating diverse results, by mapping them onto the model's variables and relationships. If results fail to fit a model, then either the model should be refined or the results should be ignored, as imprecise or irrelevant. The process of creating a model should improve an organization's internal communications by requiring experts to interpret their knowledge in measurable, mutually comprehensible terms, while encouraging them to examine interdependencies. Having computability as a goal, rather than actual computation, affords more equal status to factors that are hard to measure or weakly represented in an organization. It allows more deliberate choices regarding the investment of analytical resources (e.g., refining existing estimates versus filling gaps).

It takes advantage of the creativity of scenarios, while giving them the rigor needed for evaluation and implementation.

The opportunity for integration arises from the fact that both models and scenarios conceptualize complex systems with similar primitives: causal factors, interrelationships, uncertainties, contingent outcomes, and possible interventions. Model runs are scenarios without the life-like detail and the opportunity for coherence tests; scenarios are model runs without the commitment to precision and completeness. Neither is satisfying to those deeply rooted in the other approach. Scenarios are hopelessly ambiguous to committed modelers. Models are hopelessly ascetic to committed scenarists. At the end of the day, people will create and consume analyses that feel right to them. However, the mediated interaction proposed here could encourage more realistic models and more analytical scenarios. Managing disaster risks is too difficult an intellectual task to leave to a single analytical approach. We need all the help that we can get.

Acknowledgements The research reported here was supported by the National Science Foundation (SBR95-9521914, SES 0350493 & SES-0433152) and Canadian Defence Research and Development. We are grateful to the participants in Pandefense 1.0, Rosa Stipanovic, and Mark Huneke, as well as seminar audiences at Carnegie Mellon University, Columbia University, Harvard Business School, and the University of Washington. The views expressed are those of the authors.

References

Brown, Rex.V. (2005). Rational Choice and Judgment: Decision Analysis for the Decider. Hoboken, NJ: Wiley.

Bruine de Bruin, Wändi, Baruch Fischhoff, Larry Brilliant and Denise Caruso. (2006). "Expert Judgments of Pandemic Influenza," *Global Public Health* 1(2), 1–16.

Burns, William J. and Robert T. Clemen. (1993). "Covariance Structure Models and Influence Diagrams," *Management Science* 39, 816–834.

Casman, Elizabeth et al. (2001). "Climate Change and Cryptosporidiosis. A Qualitatitive Analysis," *Climate Change* 50, 219–249.

Clemen, Robert. (1997). Making Hard Decisions. Belmont, CA: Duxbury.

Crouch, Edmund and Richard Wilson. (1981). Risk Benefit Analysis. Cambridge: Harvard University Press.

Department of Health and Human Services. (2006). Pandemic Planning Update. Washington, DC: Department of Health and Human Services.

Dombroski, Matthew, Baruch Fischhoff and Paul Fischbeck. (in press). "Predicting Emergency Evacuation and Sheltering Behavior." Risk Analysis.

Downs, Julie S. et al. (2004). "An Interactive Video Program to Reduce Adolescent Females' STD Risk: A Randomized Controlled Trial." *Social Science and Medicine* 59, 1561–1572.

Eggers, Sara L. and Baruch Fischhoff. (2004). "A Defensible Claim? Behaviorally Realistic Evaluation Standards," *Journal of Public Policy and Marketing* 23, 14–27.

Fischhoff, Baruch. (2000). "Scientific Management of Science?" *Policy Sciences* 33, 73–87.

Fischhoff, Baruch. (2005). "Decision Research Strategies," *Health Psychology* 21(4), S9–S16.

Fischhoff, Baruch, Julie S. Downs and Wändi Bruine de Bruin. (1998). "Adolescent Vulnerability: A Framework for Behavioral Interventions," *Applied and Preventive Psychology* 7, 77–94.

Fischhoff, Baruch et al. (1981). Acceptable Risk. New York: Cambridge University Press.

Fischhoff, Baruch et al. (2003). "Evaluating the Success of Terror Risk Communication," *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 1(4), 255–258.

Fox, Maggie. (2006, May 9). "Experts Fear Bird Flu Movie May Spur Panic," Reuters.

Funtowicz, Sylvio and Jeremy Ravetz. (1990). Uncertainty and Quality in Science for Policy. Dordrecht: Kluwer.

Güvenç, Ümit. (2005). A Scenario-based Risk Assessment and Risk Communication Methodology for a New Biomedical Technology. Carnegie Mellon University. Unpublished dissertation.

Harris, Gardiner. (2005, September 14). "From Washington, a Story about a Killer Flu," New York Times.

Institute of Medicine. (2006). Reusability of Facemasks During an Influenza Pandemic. Washington, DC: National Academy Press.

Kaufmann, Arnold F., Martin I. Meltzer and George P Schmid. (1997). "The Economic Impact of a Bioterrorist Attack: Are Preventional and Post-attack Intervention Programs Justifiable?" *Emerging Infectious Disease* 3(2), 83–94.

Krayer von Krauss, Martin P. Elizabeth Casman and Mitchell J. Small. (2004). "Elicitation of Expert Judgments of Uncertainty in the Risk Assessment of Herbicide-tolerant Oilseed Crops," *Risk Analysis* 24, 1515–1527.

Krishnamurti, Tamar P., Sara L. Eggers and Baruch Fischhoff. (2006). "A Behavioral Decision Approach to Regulatory Law: An Application to Teens' Contraceptive Decision-making Competence," Under editorial review.

Levi, Michael A. and Henry C. Kelly. (2002, November). "Weapons of Mass Disruption," *Scientific American* 77–81.

Morgan, M.G. et al. (2001). Risk Communication: The Mental Models Approach. New York: Cambridge University Press.

Morgan. M. Granger and Max Henrion. (1990). Uncertainty. New York. Cambridge University Press.

Morgan, M. Granger and David W. Keith. (1995). "Subjective Judgments by Climate Experts," *Environmental Science and Technology* 29, 468A–476A.

National Research Council. (1996). Understanding Risk. Washington: National Academy Press.

Salaam-Blyther, Tiaji and Emma Chanlett-Avery. (2006). US and International Responses to the Global Spread of Avian Flu. Washington, DC: Congressional Research Service.

Schwartz, Peter. (1991). The Art of the Long View. New York: Doubleday.

Viscusi, W. Kip and Joseph E. Aldy. (2003). "Value of Statistical Life: A Critical Review of Market Estimates Throughout the World," *Journal of Risk and Uncertainty* 21, 5–76.

Wessely, Simon. (2005). "Don't Panic! Short and Long-term Psychological Reactions to the New Terrorism," *Journal of Mental Health* 14, 1–6.

WHO. (2005). Mental Health of Populations Exposed to Biological and Chemical Weapons. Geneva: World Health Organization.

Zamiska, Nicholas. (2006, January 13). "Antiflu-drug Handout in China Alarms Health Experts," *Wall Street Journal*, p. A9.