The Journal of Psychology and Financial Markets
2002, Vol. 3, No. 2, 102–111

Risk Perception, Risk Communication, Risk Taking

Baruch Fischhoff

Baruch Fischhoff: I'm with the Departments of Engineering and Public Policy and of Social and Decision Sciences at Carnegie Mellon University. What I'm going to talk about is research done over the last few years on risk perception and communication and its implications for investment decision making. Last night's talk on neuroscience mentioned George Wolford, at Dartmouth. Some years ago, he and I were at the Medical Research Council Applied Psychology in Cambridge, was along with a neuroscientist named Tony Marcel. He made the following observation regarding when we make progress in understanding the brain. He said, "We made a lot of progress after World War I because we had a lot of brain-injured people to study. However, that petered out because we didn't have the psychological tasks needed to understand their behavior. We made more progress after World War II because we had more people with brain injuries. By that time, psychology had advanced a lot, primarily promoted by the war effort, giving us tasks to understand the effects of the injuries. However, we were limited by our understanding of the anatomy out of those injuries. In the '70s it all came together. We had good tasks, good internal understanding and it took off."

That is the spirit in which I'm going to offer the sort of research that we do, most of which is not directly on finance. However, it is another set of tasks, used to study risk taking in other domains, along with some results and some small-scale theories that might be useful for understanding financial risk taking. The research that I'm going to be talking about concerns public responses to health, safety, and environmental hazards. Dealing with such risks (and perhaps the ones posed by financial markets) requires both quantitative and qualitative understanding. People must understand how big the risks and benefits are (including the associated uncertainties), as well as how these risks are created and controlled. That qualitative understanding reflects people's folk psychology regarding the systems that they are confronting.

In most of the examples I will be presenting, there's somebody responsible for creating and controlling these risks. So, if you're analyzing the industry with that responsibility, it might be of some use to you, to know what their exposures are and how the public views them. Even without that substantive interest, you might find value in the kinds of results that we have found and the methods available for other applications. This kind of psychology, like a lot of the behavioral decision research, begins with a normative analysis, putatively specifying what people ought to be doing in a situation. It is followed by a descriptive analysis of what it is that they're actually doing and a prescriptive analysis of what we can do to close the gap and help people to make better decisions.

In terms of the epistemology of this kind of research, the challenge that we face, in trying to do systematic science regarding these complicated problems is somehow matching the real world with the sort of situations that we can study in a controlled way. So our challenge is identifying the critical variables in complex situations that have sufficient generality to merit being studied in isolation.

Allan Baddeley wrote a paper about 20 years ago, identifying two aspects of a research program that sought to make statements about the real world. One is *applied basic* research: taking what we have already got and making it useful to somebody, in the process testing whether we understand our phenomena well enough to identify their implications for specific situations and anticipate how well they generalize.

The second essential component is *basic applied* research: identifying basic research topics that arise from the applications, so that we free ourselves from the trap of unduly focusing on the endogenous problems that occupy many research fields, entertaining ourselves on variations of the things we have always done. Working with practitioners is one way to free ourselves of that trap, as they will force us to address the issues that we are missing.

I will start with some research on people's judgment of quantitative risk. Some years ago, I

Baruch Fischhoff is University Professor of Engineering and Public Policy and Social and Decision Sciences at Carnegie Mellon University. He is also Director of the University's Center for Integrated Study of the Human Dimensions of Global Change.

Requests for reprints should be sent to: Baruch Fischhoff, Carnegie Mellon University, 5000 Forbes Avenue, Porter Hall 219E, Pittsburgh, PA 15213. Email: baruch@andrew.cmu.edu

was a reviewer of a book by Kip Viscusi on smoking risk perception. It came back to my attention when Paul Slovic and I (and some other people) were asked to work for the State of Massachusetts in their suit against Philip Morris. Kip had taken some data collected for the industry in which they asked people to estimate probability that a smoker would die of lung cancer. The industry lawyers claimed that the study showed people overestimating the risk, hence needing no additional warning. However, if you look at the distribution, you see that it is bimodal. There is a lump of people who think the probability is roughly 10 to 30 percent, and then there's another lump that think the probability is about 50 percent.

About the time I was working on these data, I came across a table from a study by Chuck Manski, a labor economist at Northwestern. He had responses from a random sample of US residents answering questions like, what's the probability you'll have no health insurance in the next year, you'll be the victim of burglary in the next year and you'll lose your job in the next year. There, too, we find bimodal distributions, some things in the 10 to 20 percent range and another big lump at 50 percent.

One possible explanation is that these are heterogeneous populations. Some of them live in very high crime, high job-loss areas, etc., with similar risk distributions in each case. But it does look a little peculiar.

A study by Black et al. at Dartmouth elicited women's judgments of their probabilities of getting breast cancer, and the distribution looks very similar to Viscusi's, 10 to 20 percent is one node, with a secondary node at 50 percent. At that time, we were interviewing a bunch of teenagers about their perceptions of risk. When I presented these results at a research meeting, a graduate student said, "You know, when I ask kids to give me probabilities, they often say '50/50.' I've always recorded that as 50/50. Maybe that's not what they mean. Maybe this is an inability to answer a question or an unwillingness to answer it. It's sort of epidemic uncertainty. If they don't know what to say, and we want a number, they'll give us a number."

Eventually she, Wändi Bruine de Bruin, did a dissertation on this topic. We began by going back into our archive of interviews and pulling out every reference to the word "50." We found questions and answers like: "How well do you think that it would protect a person from getting HIV, if they used a new needle?" "I wouldn't say hundred percent, I would say almost definite, but there is always risk involved. I think with that kind of lifestyle you are going to eventually run into that, so I say a 50/50 chance."

Occasionally, they say 50/50 and they mean it. For example, we pulled up one with a question regarding the probability of getting addicted if you use crack cocaine: "Oh, it's 50/50 because my mother tried it and she did and my sister tried it but she didn't."

So this person had a problem of generalizing from small samples (although that might be the least of her problems). Such results pose an interesting situation for cognitive psychologist. They suggest an emotional undercurrent to judgments about risk, of people not being able (or willing) to get their minds around some situations. By opening up our methodology a little bit, not so much as an anthropologist would, but still more than the conventional survey or experiment, we saw some things that, at the very least, show the limit to our methodology and maybe reveal something else about human behavior.

The first study we did elicited probability judgments from people snagged on Sproul Plaza at Berkeley. (You can think about generalizeability of that sample's judgments under conditions of uncertainty, as you will.) We gave them a couple of bucks for completing a short questionnaire. We asked them questions like what's the probability that an average woman will be diagnosed with breast cancer in her lifetime, with a typical open-ended probability scale, asking for a number between 0 and 100%. The second group used a linear scale, like you also find in a many studies, on which they had to make a check mark. The third group received a log-linear scale that we have often used with adolescents. It blows up the probabilities from zero to 1 percent to make it easier to give small probabilities. We have found that kids and juveniles, given a warm-up question, have no obvious trouble using this sort of log-linear scale. Their numbers look like those of adults.

Here is the distribution of probability judgments for a woman getting breast cancer. People given the blank, open-ended response mode have a big lump of 50's. You give a linear scale; the lump of 50's just about disappears. With the log-linear scale, it's even smaller. And you can get this pattern pretty routinely. If you get a 50 blip at all, then you can pretty much eliminate by giving a marked scale.

Here is the distribution with another question. What's the probability of AIDS transmission in one case of unprotected sex? You get a big blip there and it flattens out with the other response modes. There's a methodological question regarding which is the best way to do this. For a while we thought, well, the best way is to give people a

scale. The more we thought about it, the more we started to feel, maybe people really don't know how to think about some of these questions. Maybe we would be suppressing a signal that they're trying to give us, if we used a structured response mode.

Part of Wändi's dissertation involved trying to figure out when people want to express such epistemic uncertainty. It seems to be more common with judgments regarding threatening events, which you can observe when you allow people to express themselves in this qualitative way. Wändi went back home and replicated this study in her high school outside of Amsterdam, where her parents grow orchids. She found that the blip disappeared because there's no 50/50 expression in Dutch. Conceivably, the Dutch don't feel this sort of uncertainty, but most likely it goes someplace else in their thinking.

Yesterday a number of the speakers talked about what is it that people feel when they don't really understand the situations facing them, when they must make predictions and choices in a world that they, on some level, don't understand or don't want to think about.

You can also look at these questions in terms of the accuracy of their estimates. It turns out that in situations like the smoking, one where people were giving too high estimates, you take away those 50's and you get estimates that are closer to being accurate.

We have some physician colleagues at Dartmouth, Lisa Schwartz and Steve Woloshin, who are interested in patients' numeracy. They make a strong case that you don't lose anything by using a version of the log-linear scale that places a little microscope over the very small probabilities, which helps people to understand that those are really small probabilities. We've also collaborated with a grad student named Neil Stiber, who's been eliciting probabilities from groundwater contamination experts regarding the outcomes of different natural remediation processes. He had used an open-ended response mode (prior to knowing about our work). We came up with an analysis, showing an excess of 50's even among these experts.

So it seems like even experts, when they don't know quite what to say, will say 50/50 because it's something to say and it looks like an appropriate answer. With Chuck Mansi's help we got a bunch of expectations questions on the National Longitudinal Study of Youth, asking teens to give open-ended probabilities. We had lost a battle to have a ratings scale, but did get the chance to look for 50/50 effects in a very large sample. We had 20 questions, and 3,800 15- to 16-year-olds. This is as good a sample of kids as you're going to find, over-sampled for African-Americans and Hispanics. We had a warm-up question: What's the probability you'll eat pizza in the next year? The teens provided very high probability judgments. This question was designed to get respondents accustomed to the scale and to see whether they would give us extreme values.

The second question was, what's your probability of getting the flu in the next year? We got a much flatter distribution, and a spike at 50. We asked some questions that Chuck can use in studying human capital formation, in the sense of when do kids invest in themselves, such as what's your possibility of being in school in the next year? Here, we found high estimates and optimistic ones, with a little bit of a 50 blip. Another question was what was the probability of having a high school diploma by 20? Again, estimates that were not unreasonable, but probably optimistic.

We asked the young women, what's the probability you'll get pregnant in the next year? The mean answer is 4 percent, which as far as anybody knows is pretty close to the truth. We asked the guys what was the probability that you'll get somebody pregnant in the next year? The mean answer was 8 percent. Since most young women's pregnancies are from older men, some of these guys are dreaming. We believe that this reflects their judgments, rather than problems with the method. We generally found fairly optimistic to realistic answers.

Here's another example: What's the probability of being in jail by age 20? This is not too far off and it's sensibly correlated with questions on other parts of the survey, describing their behavior and neighborhoods. The final two questions on the scale ask what's the probability you're going to die in the next year, and by age 20? Here, we got very high probabilities—if you believe that the rating scale allows them to express their beliefs. A minority of teens give zero or 1 percent, even though the actual annual probability is 0.08%. So they're giving you non-negligible probabilities and a big spike at 50. The median probability of dying in the next year for these kids is about 20 percent, and it almost doesn't differ by gender and race. The white guys get slightly lower probabilities, but otherwise all the kids are up there. If you take a median, it goes down to 10 percent (because the 50s pull up the mean). Many teens give a higher probability to graduating college by 30 than they do of living to 16 or 17.

So what's going on here? Yesterday, Don MacGregor was talking about the images that people have about the future. It seems like there's some instability in the images that these kids have of their

own futures. It seems as though they can think of themselves in various future states, regarding school and work. However, they can't integrate the pieces into a coherent view of the future.

We also asked these kids what was the probability that they would die by age 20, which was an extra three-and-a-half years on average, and it was about the same probability as for the next year. In some fundamental sense, either they don't know how to or they don't want to think about these possibilities. However if this reflects something about how they're thinking about the future and about how investors think about the future, maybe there's something here regarding how people face situations that they find imponderable. These are beliefs that you don't get to hear unless you allow people to express them.

Voice: Did the question include a category of don't know or refuse to answer?

Baruch Fischhoff: We didn't on this study. However we've done that in an experimental one, where we've added, "I don't know" or "I have absolutely no idea." Doing so somewhat reduced the rate of 50s, without markedly changing the rest of the response distribution. One could also encounter reluctance to respond when respondents are uncertain what the questions mean. We worked very hard on those questions to get ones for which we could find statistical accuracy, which required us to eliminate ambiguity. The only one where we couldn't get a statistical estimate was how many people eat pizza per year. We know how many acres of pizza are eaten a year, but not that.

Arnold Wood: Are these solitary experiments? In other words, does just one person do it at one time?

Baruch Fischhoff: Yes.

Arnold Wood: You don't know the answers other people are giving?

Baruch Fischhoff: All this is individual.

Arnold Wood: It's not verbal, it's written down?

Baruch Fischhoff: The short transcripts that I showed you are from one-on-one interviews. These other results are from surveys. If you want to evaluate the quality of people's beliefs, you have to understand how they are interpreting their questions. Here's a question from the National Center for Health Statistics survey, one of our big surveys. It asks how likely do you think it is that a person will get the AIDS virus from sharing plates, forks or glasses with someone who has AIDS?

If you wanted to evaluate someone else's answer (or provide one of your own), you might want to know how much sharing. Is this sharing on a single occasion, or several occasions? And you might also want to know how intensive that sharing is. We asked a bunch of undergraduates (at Yale) to answer the question and then to tell us how they had interpreted the question that they had answered. We had a fairly flat distribution of interpretations of the dose, or amount of exposure, and a fair degree of agreement about the kind, or intensity, of that exposure. But even that would only be useful to you as an investigator in interpreting responses if you happened to know the culinary norms of Yale undergraduates, or of any other group, on the basis of which they interpreted such ambiguous questions.

If you ask people to think aloud as they're answering these questions, you get answers like this one, in response to a deliberately ambiguous question, similar to those found on many surveys, "What is the probability that a person who smokes will get cancer?" Here is a teen, I think it's a kid in juvenile detention, he said: "It's about 50/50 because it depends if it's in your family. There are a lot of questions, you might just quit smoking. You might quit, you know, start striking for a couple of weeks and so on," and it goes on, "but it's about a 50/50 chance." At that point we didn't realize that 50/50 chance had some sort of special status.

If you talk to these kids, they're much more sophisticated than the investigators. That is, here you've gotten an unqualified question and they're trying to specify it in a way that allows them to give a more precise answer. It turns out that if you do this, you find certain kinds of regularities. We posed nine such questions to 60 kids. We found that in most cases, they wanted to know how much of the dose was involved.

For example, how much marijuana before you get addicted, how much alcohol before you get addicted. There were two questions, regarding sexuality-related risks. We asked them what's the probability of getting AIDS if you have sex? They didn't feel the need to know how much sex you would have. And the same thing for pregnancy. In part, this insensitivity reflects teens' folk theories about sexuality. But it's also partly the case that kids—like everybody else (there's studies on this going back to the 1950's)—have difficulty seeing how quickly risks mount up through repeated exposures. You have sex, either there's a pregnancy or not, unlike smoking, where you can imagine the substance and damage building up within your system.

So with things that either happen or don't happen, people tend to think tactically (in terms of single exposures), rather than strategically. If you give them a probability and ask them to accumulate that probability through multiple exposures, they tend to under-accumulate the risks. Such short-term thinking, and failure to appreciate cumulative effects, may be close to investment-like behavior.

Let me shift to the next topic. Imagine that you have a duty to inform people about a set of risks. You've done your best job of figuring out what they currently believe those risks to be. You want to use their time well by giving them the information most important to refining those beliefs. You have a small window of opportunity. You want to give them the information that's most relevant that they don't know that's most relevant to their choice.

I had a grad student named Jon Merz, who is currently in the bioethics group at Penn, who did his dissertation trying to come up with an analytical procedure for deciding how it is you prioritize information. John took a medical procedure, carotid endarterechtomy. In it, they scrape out the carotid artery, for people vulnerable to stroke. If you're a candidate for the operation and you get through the surgery, then you're better off. But you might not get through the surgery, and there are lots of things that can go wrong. Here's a list beginning with the fact that you could die from the surgery, per se. It could precipitate a stroke, although, with a lower probability than having a stroke without the surgery.

The details of what Jon found are in one of the two papers that I handed out. Jon had been a nuclear power plant risk analyst and a patent attorney before he came to study with us. Drawing on those skills, he did a quantitative risk analysis, in which he summarized the literature, regarding the magnitude of the risks from these different sources. Then he estimated the percentage of patients who would change their decision, if you told them about each of these different risks.

I won't go into the gory details. However, if you believe his procedure, then if you told people what the probability of death was from the surgery, about 15 percent of patients would be close enough to the threshold that they would back away, perhaps because the benefit was not that great for them or perhaps because they are particularly sensitive to the short-term risks (of surgery) as opposed to long-term risks (of living with the condition). If you tell them about the possibility of stroke and neurological deficit, another 5 percent of patients would say, "No, this isn't for me." You tell them that you could have nerve paralysis; another 3 percent decline the surgery.

As it turns out, there are three risks of this surgery that is absolutely essential to know. You wouldn't want to hide anything. However, if you had a duty to inform, you'd want to make certain that people understood those three risks. If you think about what it takes to understand those risks, then every risk has a probability and an event, both of which need to be understood. It turns out that the probability of death, the probability of stroke and the probability of nerve paralysis, they're all large enough that they're not hard probabilities to explain. They're not one in 56,000.

So you can focus your efforts on explaining the events. The first event, death, is probably impossible to understand, but there's not a whole lot you can do to help somebody. Stroke, these are often people who have had a stroke, but you could explain a little bit more about that. As a result, you might end up spending a lot of your time explaining what this nerve paralysis is all about. It's the third thing on the list, but perhaps the least known.

These are hypothetical populations modeled after the actual populations. What John did is to create a hypothetical population of people who would be candidates for the procedure—assuming that money was no object and there were no side effects. Then he said imagine that these people were rational decision-makers and you told them about each side effect, how much of a difference would that make to them?

In this situation, there's no reason why anybody would know about the surgery. As a result, you can assume that they are starting from ground zero—they know nothing. As a result, you didn't need to do the psychological assessment of current knowledge, in order to prioritize this information.

In the domain of finance, you could apply the same perspective, thinking about what you would want to say if you had a duty to ensure that people understood the risks of an investment. Initially, you would analytically determine what it is that they need to understand. Empirically, you could determine how much they know already. Practically, you could develop a communication to close the gap, then evaluate the your success. Similarly, you could evaluate the quality of the evidence that other people present in their communications to investors.

Here's another example of how to design a warning, also drawn from a grad student's dissertation, Donna Riley who came with a chemical engineering degree from Princeton. There was a time, about four years ago, when the Consumer Products Safety Commission, the EPA, and part of the chemical industry were trying to voluntarily deregulate the management of hazardous chemicals, one of which is methylene chloride. It's a very nasty, but also very effective solvent, one of whose uses is as paint stripper.

The regulatory question was could you get people to voluntarily reduce their exposure to a safe level, so that you wouldn't have to pull it from the market? This was a behavioral question, upon which a large investment hinged. What Donna did was take the specifics of the situation seriously,

and create a behaviorally realistic exposure assessment. She started with a quantitative risk exposure model, then tried to figure out what the diffusion patterns for methylene chloride are—it's a pretty heavy gas—in a typical basement workplace. Then she did some interviews to find out what people thought about the chemical and how they used it. She then incorporated these behavioral possibilities into an otherwise conventional risk assessment.

These analyses allowed her to identify the most effective ways to reduce risks, and to evaluate the effectiveness of communications designed to achieve that goal. Methylene chloride is made by a few companies then sold to many formulators who produce the retail product. There are probably analogs within the security industry, whereby somebody generates a product that many people sell, putting their spin on it, their own advertising, their own wrapper, formulating it a little bit differently. You could ask how well do their warning labels inform investors about their risks.

With methylene chloride, there are two risks. One is absorption, where the main risk is that the chemical is a probable carcinogen. The second risk is that it also gives off carbon monoxide so people have heart attacks from using it in poorly ventilated places. There are lots of things that you might recommend to people. Donna tried to estimate the risk reduction from each of these things—assuming either that people religiously followed what you told them or that people realistically followed what you told them.

Some things people won't do because it's hard to get the job done, even though they know it's the right thing. Some things they will do, but they'll do it wrong. So if you tell people to ventilate the room, some people will have the fan blowing in rather than blowing out. Donna estimated how much protection you could get from each recommendation. If anybody is interested, open the window; it's usually the single best thing you can do. The second best thing is getting goggles because there are perhaps 30,000 emergency room admissions a year with methylene chloride in the eyes.

In order to evaluate how well current producers had fulfilled their duty to inform, Donna estimated much risk people would have if they looked at the labels of each of several different products, making different assumptions about how they read the label. One analysis assumed that they read the first five things on the label and followed them religiously. It turned out that doing so would have no effect on the risk level with many products. With one, though, it would have a significant risk reduction; the relevant information was right there.

On the other hand, if you only read the emphasized text that same label wouldn't help you. This is very detailed analysis on a single product. But this is representative of our way of trying to figure out what information is critical and how good a job are we doing with informing people—considering both the risk of the situation and the psychology of the individual.

The regulatory impact of this exercise proved to be smaller than it might have been. It turned out that the EPA people were called off on something else. The Consumer Products Safety Commission people were very interested in this as a different way of doing business, enough so that they said let's check with our commissioners. Unfortunately, at that time, they didn't have a quorum of commissioners, so they couldn't do anything. The industry people remained interested because they represented manufacturers, who were afraid of the risks created for them by the retailers, with their inadequate labeling.

From a regulatory perspective, by having a consistent standard, you would have a level playing field in a situation where nobody wanted to be the first to bring risks to consumers' attention. With methylene chloride you want to know what to do. However, you are less likely to do it unless it fits with your theory of how fans works, rooms ventilate and so on. If so, then you need some theory of the process and not just a bottom line recommendation, or statement of risk levels.

Sometimes you need a bigger picture when you're trying to get a realistic feeling for why they tell you the risk is of this magnitude. My next topic concerns how one creates such bigger pictures. Recently, we had a project with the American Waterworks Research Foundation, who wanted us to design the perfect boil water notice for people who had cryptosporidium in their water supply. This is a water-borne parasite. It was the source of an outbreak in Milwaukee a few years ago, where about a hundred people died and 404,000 people got sick. We approached this task in the same spirit as the methylene chloride study, first analyzing the risks, then seeing what people currently believe, then determining what they most need to know in order to take care of themselves, should a problem arise.

Half the people we interviewed were HIV-positive. Among people who have AIDS, there is a 60 to 80 fatality rate with cryptosporidium, because it goes directly to the liver and just explodes. If you have a compromised liver, it's very dangerous. So we interviewed people who are HIV-positive, a few of whom had AIDS, as well as people in certain areas that had had crypto outbreaks to find out what they believed about the problem and possible

responses, including how they boiled water. Some people were very focused on the task, others defined it very broadly, wanting to know about the animal feed lots that increase the risk. Others wanted to know, who are these people giving me the boil water notice? Whom are they protecting? Why are we doing boil water notices instead of better land-use planning? What if I hear about the problem from TV, rather than from the health people?

So we ended up building a whole model characterizing the things that determine the risk. The model is created so that you could set the conditions for any place with, say, a particular probability of runoff or of detection. You could consider how good is the media coverage? How good is the feedback from the risk managers? Once we had the model, we could calculate the value of the perfect boil water notice, which took people from their current proficiency at boiling water to perfect proficiency (or to realistic proficiency). We could ask how much could you reduce the crypto risk?

From our interviews, we discovered that the people don't do a very good job of boiling water now. However, once we ran the numbers, we discovered that, even if they were perfect, it would make no difference to their risk. We puzzled over this result for a couple of weeks, because it meant that our communication project would be useless. Once we thought about it, we realized that the reason was that crypto is very hard to detect. It has to be cultured in feces, which is very expensive and the detection still isn't very good. You've got to get the medium just right, so it takes a week between when you collect the sample and when you get the results. After a week, the result is only good for forensic work. If you want to know what hit you, this is good a good test.

Now we have the perfect boil water notice for e.coli, whose effectiveness can be predicted with basically the same model. Because e.coli can be cultured quickly and accurately, a good notice can have a big impact. But for crypto, this whole system was relying on boiled water notices that could have no effect whatsoever. So again, if you step back and do this kind of model, sometimes you discover fundamental aspects of your risk, which were not obvious, however easy it is to explain them, once revealed. It also shows you something about the nature and complexity of the problem that people are dealing with. Here, the problem was so complex that a simple critical relationship was not obvious to the people in control.

Voice: I think what you just described is almost a perfect analog to what investment management organizations struggle with when they try to put a risk description in the prospectus. They're really attempting to describe something that not only do their customers not understand at all, but the investment organization, itself, may not understand very fully. And do you have any thoughts on how in, say, 1999 a mutual fund organization could have sensibly informed people about investment risk at a time when investments appeared to have none?

Baruch Fischhoff: Well, that's what I thought about the relevance of the research, but I wanted somebody knowledgeable in finance to suggest that. Let me show you how we approached a risk problem with primarily financial consequences. Ontario deregulated its energy markets starting in '97, and for reasons having to do with how the Canadians run their country, they mandated a wide-scale consultative process to ask what ordinary citizens, people in the different user communities, indeed everybody who has a stake in electricity, thought about it. The province basically divided the old Ontario Hydro into two parts, production and transmission assets. We were asked by the transmission company, now called Hydro One, to consult with their publics on how to set the rate structure for electricity distribution. This is something barely understood by people on the inside. Yet the province insisted that everyone from First Nations to lobbyists have input to the process. They were so desperate that they involved psychologists in it.

As usual, our process began by building one of these models. We called them "expert models." It's a bit of a euphemism, when, as with the crypto model, we get critical input from non-experts showing us that we've bounded the system too narrowly. Some situations are so complex and novel that even your experts don't know all that much. The approach is one that we borrowed from our engineering colleagues who did the first model like this on an acid rain problem, where they had a need to coordinate people who knew about fluidized bed combustion, scrubbers, atmospheric transport, atmospheric oxidation, groundwater transport, and so on, but couldn't talk to one another.

They came up with the formalism of this kind of influence diagram, as a way to figure out how the pieces fit together. The idea is to capture the essential issues in the system's behavior, provide a common framework and language, allow access to specialists and non-specialists, and permit reliable coding of inputs. Imagine that either an expert or a citizen says I'm really worried about how this change is going to affect jobs in Ontario. You have a way of systematically characterizing what they said and pooling it with other beliefs and research on the topic, as a step toward identifying recommendations.

This is our strategy for creating models to design communication processes and products. You

would have to tell me whether there's any relevance to the equivalent tasks with different classes of security risk, including that of investment bubbles. If asked to help, we would review the technical literature, find the appropriate representation language, draft a detailed model, test it against objective descriptions of the system, and so on. This is standard procedure with any sort of modeling. The exception here is the commitment to incorporating knowledge of whatever types are relevant to the problem—not just some classes of readily quantifiable variables.

The model that we came up with here had this basic structure. The rate structure for transmission turned out to be phenomenally complicated, and to have complex effects on the power system, the environment, the economic success of the province, and so on. Using this model as a template, we managed to interview a lot of people and hold public hearings, including special visits to dominant economic interests and First Nations. The result was a strange sort of heterogeneous peer review, which ended up picking up a bunch of problems in the initial policy proposal. Even Hydro One felt that they had learned things in terms of how this complicated system would operate. And their proposal sailed right through the OEB, which basically approved both their process and its products.

This is a process that allows you to integrate psychology with a technical understanding of the system. It provides an appraisal of what the risks are and guidance on how to structure your communications.

The final thing I'd like to talk about is how people put these pieces together in order to make a decision. Here's a project we completed a few years ago. We interviewed a bunch of adolescents, asking them to tell us about decisions in their lives. These are kids, but they had no difficulty providing a lot of decisions. Toward the end of the interview, we said pick two hard decisions that you made and one hard decision that you're trying to make and tell us about them in detail. We looked at the intuitive structure of their decisions, in terms of the formal structure of decisions.

Here is an example of how these adolescents think about important decisions in their lives. "My cousins—ordinarily, I'm real close with my cousins and everything. My cousin was having this big graduation, but my friend, she used to live here and we started preschool together, you know, and then in the seventh grade her dad—stepdad—got a job in Ohio so she had to move there. And she was in Ohio and she invited me up for a weekend. I always have so much fun when I go up there for a weekend and I wanted to go to both things so bad, like. I think I wanted to go more, like, up to Ohio, you know, to have this great time and everything, but I knew my cousin, it would be kind of rude to say, well, my friend invited me up for the weekend. My cousins were out of town—were coming and everything, and I wanted my mom to say, 'Well, you have to stay home,' so then I wouldn't have to make a decision. She said, 'I'm not going to tell you you have to stay home, you decide.' I hate it when she tells you that you have to decide. So I decided to stay home, basically, because I would feel really stupid and rude telling my cousin, well, I'm not going to be there. And I did have a really good time at her graduation party, but I was kind of thinking I could be in Ohio right now."

We had a couple of hundred decisions to analyze. One of the things we looked at was the structure of the choices that they described. For example, how many options are they thinking about? Are they thinking about one option (whether to do something), two choices (I could do this or I could do that), a finite set of choices, or in terms of "what could I do" (in the sense of designing options)?

We discovered that many kids were not even deciding whether or not to do something; they were actually deciding whether to do something that they were already doing. That seemed like a very primitive outcome structure. What we found was that most of the decisions have to do with a single option. Very few had to do with any kind of strategic planning or big picture stuff.

If you think about the research literature on people's insensitive to opportunity costs, this is quite compatible. You only have opportunity costs if you think there are opportunities. When we think about the sorts of decision that people are forced to make in the hyper-complex environments created by investment decisions, it would not be hard to imagine them failing to fill out their option space. Thus, what we see with teens might be a reflection of how even adults cope with some of their choices.

I had another teen story here. The final line is: "Like, I decided between my parents and a boy and I just picked my parents." Maybe a good place to end, but let me talk about one other study that we've done, and offer a few concluding remarks.

We have just finished a project trying to reduce the risk of sexually transmitted diseases, by addressing the sorts of decision-making processes revealed in interviews like these. We ended up developing a DVD intervention aimed at young women. We evaluated it in a randomized clinical trial. In preparation, we asked girls to talk about their decision making when dealing with guys.

What we found is that once they got into a party situation, they weren't making any decisions. We

kept trying to change our interview protocol to get them to describe choices and they just gave us scripts: You go to the party and this is what comes. You don't want to have sex, don't go to the party. So we thought what we might be able to do is help them to see that there are decisions and then to act on those possibilities. We provided an interactive video. It included a dramatized pelvic exam, designed to provide some information and model how to talk to your physician, enlisting decision-making help.

We had critical information about STDs. We had a big sister showing a little sister how to get condoms and put one on a cucumber. But the core of it was 4 choose-your-own adventure stories, where two young women dealt with guys, one was a pickup and the other was a steady. Both guys kept, like guys occasionally do, pressing them for more. We show users that in each of these fairly realistic situations, you could still think there was a choice to be made. And we show a model for how you could get through it, using a strategy cognitive rehearsal. With it, you think about how you're going to respond, in a way that exercises those mental muscles.

A hundred girls completed the course. They saw this interactive video, four times over six months. Another hundred completed a print version to see whether the medium, per se, made a difference, and a hundred viewed some existing brochures, matched for content. We have a really good school of fine arts and had some really good actors and actresses in the DVD. We found that it worked. We had a high-risk population; about 17 percent had chlamydia coming in, as well as other problems that they were treated for.

We looked at recurrence rates. We found that there was much lower recurrence with the kids who had this kind of involving intervention, where we had taken the trouble to figure out what was really important to them and trained them to act on it. They also better reported condom use, and so on. So if one thinks about risk communication and risk taking, the critical questions are what would you like to do, what do you need to know in order to make those choices, what do you know already, and what do I need to tell you in order to make you capable to decide effectively in your own best interests.

We've got a bunch of projects where we've tried to follow this strategy. These studies are high on detail, and low on sweeping generalizations about people and their behavior. We focus on detailed analysis of the help that people need to cope with the situations that they face. I think any meaningful work needs requires collaboration between behavioral and substantive specialists. I think this is the sort of research that is absolutely essential for us to know what we can say, responsibly, about behavior in financial markets.

One conclusion is that you could audit financial procedures for their behavioral realism. For example, are you really informing people about the risks that they're taking? And also are you analyzing the risk in a way that takes the full situation into consideration? It might be interesting to look at what was considered in California's energy deregulation, and the associated risk statements, relative to what was in the picture from Ontario. You could see whether they left the politics out, as Vernon [Smith] was saying, in a way that would have told you that maybe you don't want to go forward with this investment. I think that this sort of collaboration with financial experts would be a real challenge for us, extending the generalizability our methods and results to your domain.

Thank you.

Arnold Wood: Time for one or two questions.

Voice: So what you have there are you have reality and you have a fantasy system and they operate side by side. And the same thing occurs in investment decision-making. I can give you the same example from a money manager who bought a stock that was at 18 dollars, the stock went to 25. And the manager wanted to sell it out. Well, all fundamentals suggested this was going higher, why did she want to sell it out?

When she started talking about her fantasies about it, what she was operating on was the idea that when something good happens, something bad is sure to follow. And that fantasy was getting her to sell when, in fact, she shouldn't have been selling given the fundamentals. So how do we somehow integrate the unconscious fantasies of people and risk that are totally different than what we're telling them rationally?

Baruch Fischhoff: I think that's the million-dollar question. To go back to where I was starting, this is a research tradition that begins with a normative analysis. And in this, psychology has benefited tremendously from economics in that we have a sharper understanding of the nature of tasks. We have also benefited from our collaboration with substantive specialists, forcing us to be more context sensitive than psychology often is. Our strength is in the analytical and cognitive. However, we need to pursue the issues in ways that open ourselves to observing such expressions of epidemic uncertainty—of the mind blanking over.

Our goal is to expand somewhat the envelope of possible rationality and control. All decisions, like that of your client, occur in some larger context. We

think about kids making bad decisions because they're impulsive. I think they're often impulsive because they can't make good decisions. They don't know what to do, it feels hopelessly complicated, their mother won't tell them what to do, and they don't know themselves. That confusion encourages impulsive resolution of intolerably indeterminate decisions.