Risk Analysis, Vol. 17, No. 4, 1997

Distinguished Award

The Words of Risk Analysis

Stan Kaplan[1]

Received January 28, 1997; revised June 17, 1997

This paper is a transcript of a talk given to a plenary session at the 1996 Annual Meeting of the Society for Risk Analysis. Its purpose is to contribute toward a single, uniformly understood language for the risk analysis community.

1. INTRODUCTION

Upon being informed that I would be expected to speak today I asked myself what I could say, in half an hour, that would be useful. Upon doing that I got an immediate answer. A voice in my head said, "Talk about the words." My first thought was that this would take more than one-half hour. Second thought was: I can squeeze it, and I'd like to do it. Besides, since I don't hear voices all that often, maybe I should pay attention. So here goes.

The words of risk analysis have been, and continue to be a problem. Many of you here remember that when our Society for Risk Analysis was brand new, one of the first things it did was to establish a committee to define the word "risk." This committee labored for 4 years and then gave up, saying in its final report, that maybe it's better not to define risk. Let each author define it in his own way, only please each should explain clearly what way that is.

2. PROBABILITY

Moreover, the discipline of risk analysis, as you know, is heavily entwined with the subject of Probability. In that subject the semantic confusion is legendary (see, for tip of the iceberg, Refs. 1–4). It has often been called a Tower of Babel. People have argued about the meaning of the word "probability" for at least hundreds of years, maybe thousands. So bitter, and fervent, have the battles been between the contending schools of thought, that they've often been likened to religious wars. And this situation continues to the present time.

To help sort out this situation, I've prepared Table I. There are three major meanings of probability, and several submeanings. The first is the statistician's meaning, which I call "frequency" or "fraction." This refers to the outcome of a repetitive experiment of some kind, like flipping coins. It includes also the idea of population variability. Such a number is called an "objective" probability because it exists in the real world and is in principle measurable by actually doing the experiment.

In contrast, the Bayesian meaning of probability, which is degree of confidence or degree of certainty, does not exist in the real world, it exists only in our heads. For that reason it's often called "subjective" probability. But that usage is misleading. It's a misunderstanding, and it has caused enormous confusion and controversy. I will come back to this point later on.

Third is the mathematician's meaning. To him, a probability curve is a mathematical abstraction. He is interested in the formal properties of such curves independent of their interpretation.

Also, there are a bunch of more recent theories that have been invented to fix alleged deficiencies in the traditional ideas. There's Possibility Theory, Dempster/Shaefer Theory of Evidence, Higher Order Probability Theory, etc. Notable among these, and currently in vogue, are the fuzzy theories (e.g., Ref. 5), which attempt to encompass, in addition to the traditional meanings, the notions of ambiguity, vagueness,

[1] Bayesian Systems Inc., and PLG, Inc.

0272-4332/97/0800-0407$12.50/1

Table I. Linguistic Chaos

Traditional meanings of "probability"			New theories		
Statistician's (frequency, fraction)	Bayesian (probability)	Mathematical (probability	Fuzzy theories (fuzziness)	Possibility theory	Dempster Shafer (relief)
Random	Belief	Formal probability	Ambiguity		
Variability	"Personal" probability	"Axiomatic" probability	Unclarity		
"Aleatory" probability	Subjective probability		Vagueness		
"Objective" probability	Uncertainty		Ill-defined		
Stochastic ontological	Confidence				
"In the world" probability	Epistemic probability				
Reliability	Forensic probability				
Chance	Plausibility				
Risk	Credibility				
	"Evidence Based" probability				

	Notation	Example
❖ What can happen? (What can go wrong?)	(s_i)	Fire/explosion.
❖ How likely is it? (What is its frequency/probability?)	(ℓ_i)	.01%
❖ What are the consequences? (What is the damage?)	(x_i)	$100,000 Two injuries. Environmental problems. Embarrassment, reputation.

Fig. 1. The three risk questions.

lack of definition, and also of paradoxes, such as the famous one about the barber who shaves those and only those who do not shave themselves. Does this barber shave himself? Well, if he does, he doesn't, and if he doesn't, he does. That's the paradox.

In fuzzy theory the answer to this question is "0.5." From the traditional side of the line the answer is, "It's a foolish question." Ask a foolish question, and you get a foolish answer. No such barber exists, nor could he. So why ask about his shaving habits?

3. TWO COMMUNICATION THEOREMS

Incidentally, I don't think there are any deficiencies in the traditional ideas. But that's another whole story. The point for now is that, with all these meanings and viewpoints kicking around, it's no surprise that there have been communication problems, big time. After struggling with these problems for a number of years, I was moved to formulate two theorems on communication, as follows:

- Theorem 1: 50% of the problems in the world result from people using the same words with different meanings.
- Theorem 2: The other 50% comes from people using different words with the same meaning.

These are actually very useful. I have them hung on the wall in my office. When an argument gets going around the table, I'm often able to point and say, "This is a case of Theorem 1 (or 2)." It's amazing how that drains the emotion out of the argument.

4. DEFINITION OF RISK

Coming back now to the definition of risk. In the first issue of the journal of the Society for Risk Analysis, John Garrick and I published a paper entitled, "On the Quantitative Definition of Risk."[6] We argued (see Fig. 1), that when one asks, "What is the risk?" one is really asking three questions: What can happen? How likely is that to happen? If it does happen, what are the consequences? The answer to the first we called a scenario, and we denoted the ith scenario by S_i. L_i then denotes the likelihood and X_i the consequences of the *i*th scenario.

So, as in Fig. 2, the triplet (S_i, L_i, X_i) constitutes "an" answer to the three questions. If we put curly brackets around it, which is mathspeak for "set of," we obtain a set of answers. We then added a "*c*" after the bracket to mean "complete" and thus to emphasize that we really want to know all the possible scenarios, or at

- What can happen?
- How likely is that?
- What are the consequences?

"An" Answer $< s_i, \ell_i, x_i >$

Set of Answers $\{ < s_i, \ell_i, x_i > \}$

Complete Set $\{ < s_i, \ell_i, x_i > \}_c$

$$R = \{ < s_i, \ell_i, x_i > \}_c$$

Include s_0 = "As-Planned Scenario"

Fig. 2. Quantitative definition of risk.

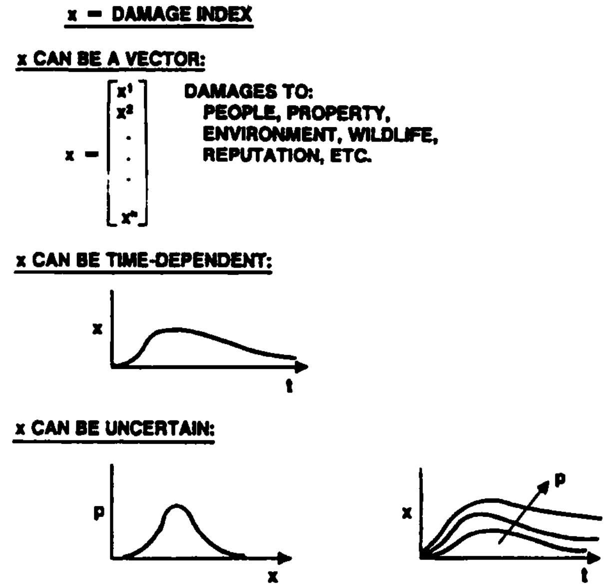

Fig. 3. Expressing the idea of damage.

least all the important ones. So we defined risk, then, as the complete set of triplets.

Notice that defined in this way, risk is not a number, nor is it a curve, nor a vector, etc. None of these mathematical concepts is "big" enough in general to capture the idea of risk. But the set of triplets, we find, is always big enough, and if we start out with that, it always gets us on the right track.

Since we called the risk scenarios S_i, we found it convenient to use S_0 to denote the "as planned" or "success" scenario. This turns out to be a very useful piece of language because if you're starting a risk assessment, the first thing to do is to write down a very clear description of S_0.

Turning to the damage index, x, we note, in Fig. 3, that this could be a vector or multicomponent quantity, that it could be time dependent, that it could be uncertain, and if so, one should express this uncertainty by giving a probability curve against the possible magnitudes of x.

Turning to the likelihood term, we noted that there are three formats with which to capture and quantify, the intuitive idea of "likelihood."

Format 1. (Frequency) This applies when we have a repetitive situation, and we ask, "How frequently does scenario i occur?" In this case the likelihood is expressed as a frequency $l_i = \phi_i$ and risk becomes $R = \{\langle S_i, \phi_i, X_i\rangle\}_c$

Format 2. (Probability) When it is a "one shot" situation, like a mission to Mars, we want to quantify then our degree of confidence that the mission will succeed. In this case likelihood is expressed as a probability $l_i = \rho_i$ and the triplets become $R = \{\langle S_i, \rho_i, X_i\rangle\}_c$

Format 3. (Probability of Frequency) The third format applies when we have a repetitive situation, or can imagine one as a thought experiment, so that the frequency exists, but since we haven't done the experiment we are uncertain about what that frequency would be. We therefore express our state of knowledge about that frequency with a probability curve. We call this the "Probability of Frequency" format. $l_i = \rho_i(\phi_i)$; $R = \{\langle S_i, \rho_i(\phi_i), X_i\rangle\}_c$

Of the three formats, we found that this third one is the most general and by far the most powerful and useful idea, and so we adopted it into the triplets and thus arrived at the full blown definition of risk as follows:

$$R = \left\{ \langle S_i, p_i(\varphi_i), p_i(X_i)\rangle \right\}_c \qquad (1)$$

Scenario ↗ ↑ Likelihood ↖ Consequence

From this definition, given the curves for frequency and damage, we can draw the various risk diagrams, such as the so called complementary cumulative, shown in Fig. 4.

5. DOSE RESPONSE

John and I thought that the definition in Eq. (1) was totally general, and applied to all kinds of risk. Indeed

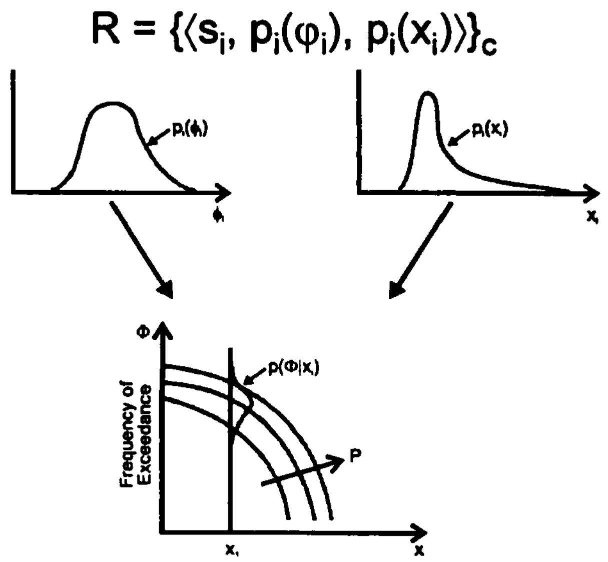

Fig. 4. Graphical portrayal of risk.

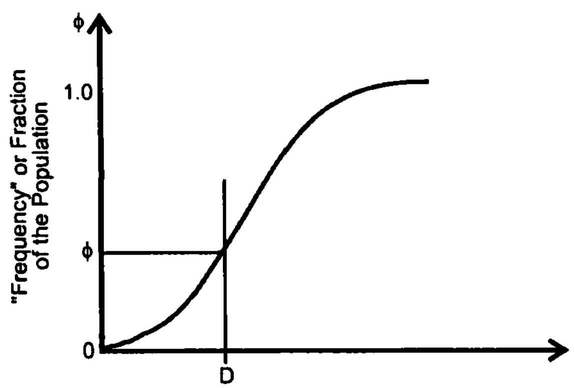

Fig. 5. Dose response curve, frequency format.

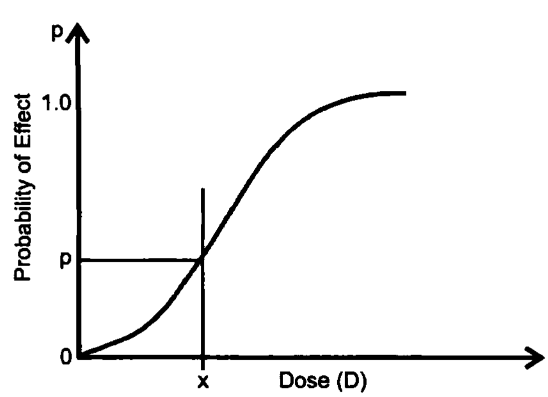

Fig. 6. Dose response curve probability format.

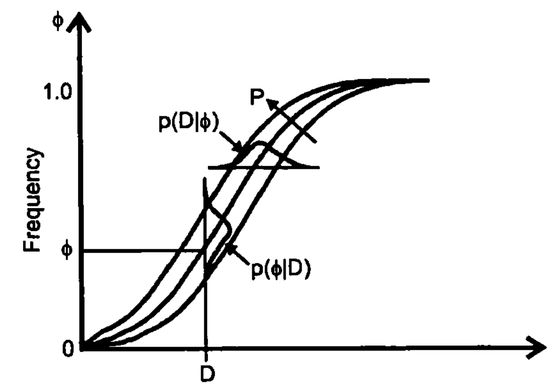

Fig. 7. Dose response curve "probability of frequency" format.

we applied it to engineering risk, investment risk, risk from agricultural pests, programmatic risk, strategic risk, environmental risk, etc.

So we thought this definition would fit everybody's needs. However, many people in our Society use a different definition. When they say "risk" they have in mind a dose response curve (Fig. 5). So the question comes up: Is this really a different definition?

There's a way of looking at it which brings it within the triplets definition. We get this curve, of course, by giving groups of lab animals doses *D*, and plotting the fraction that get sick. If we define the success scenario S_0 as the animal remains healthy, and S_1 that it gets sick, then we can see this as fitting the triplets definition, using format 1, frequency, for expressing likelihood.

If we were dosing a single animal we would plot *p*, our degree of confidence that it will get sick (Fig. 6). This fits our definition using format two for likelihood.

Suppose we imagine giving the dose to an entire population of animals (see Fig. 7). Since we haven't done this experiment we don't know what fraction would get sick. But we know something about that, so we express what we know as a probability curve against that fraction.

Connecting the percentiles we then get a "band" of dose response curves expressing our state of knowledge about the outcome of the contemplated experiment. This can be viewed as the dose response curve in "probability of frequency" format.

In connection with this curve, much controversy centers on what it looks like in the low dose range. Is it concave, convex, linear? If we plot it on a log scale (see Fig. 8), the question changes to: What's the probability curve at a low dose? That curve may be very broad on the downside, because to distinguish between an illness fraction of 10^{-3} and 10^{-4} we'd have to do an experiment

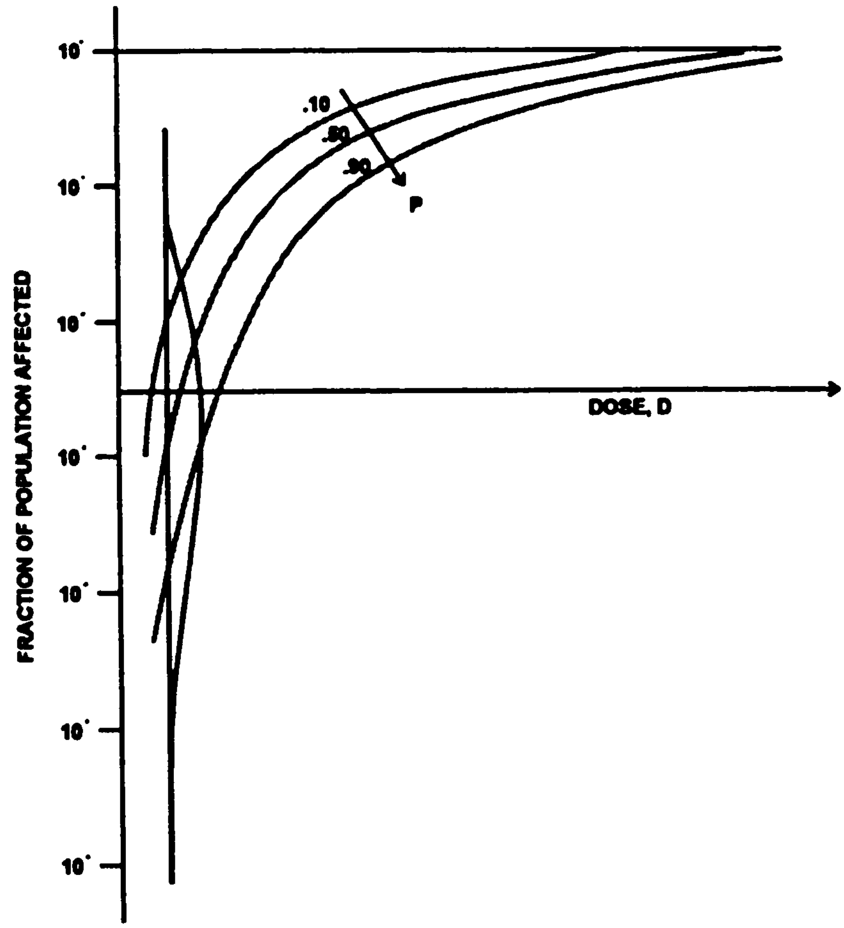

Fig. 8. Dose response curve in low range.

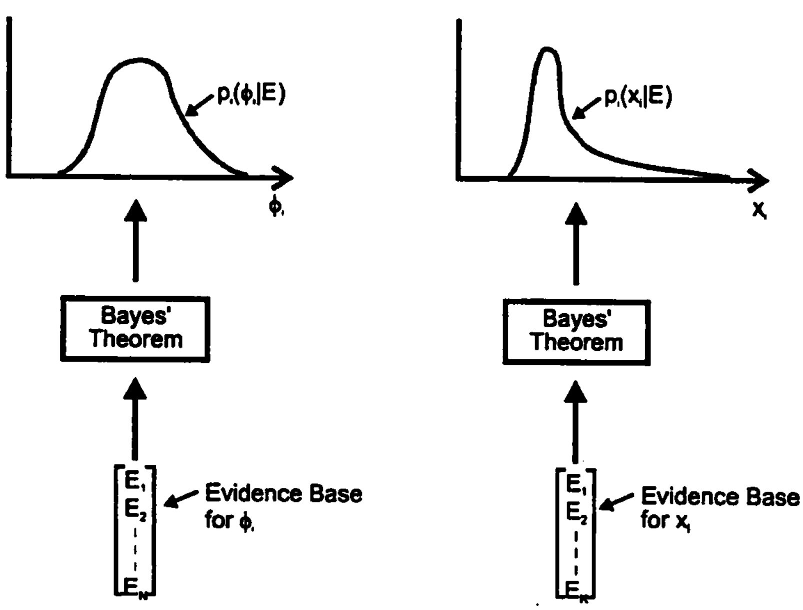

Fig. 9. "Evidence-based" approach.

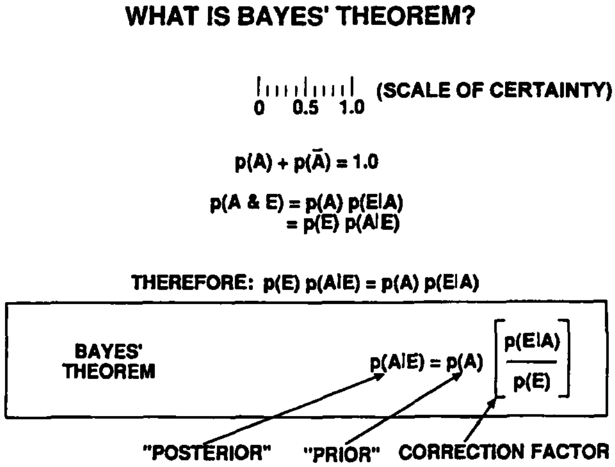

Fig. 10. Derivation of Bayes' theorem.

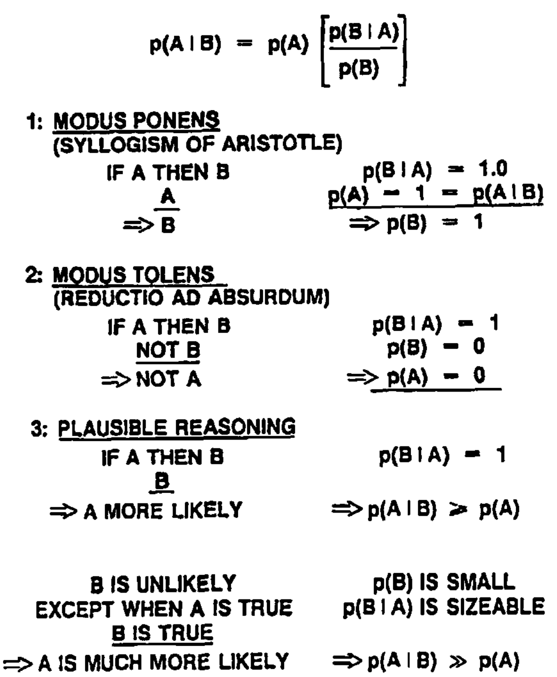

Fig. 11. The fundamental principles of logic seen as special cases of Bayes' theorem.

with thousands of mice. A hundred is not enough to tell us where the true fraction is in that range. It is enough, however, to tell us where it is not. If none of them get sick, for example, we have high confidence that the true frequency is not in the 10^{-1} range. But it's not enough to tell us where it is. So unless we have some other evidence, this is the curve we have to carry forward into our decision process.

6. BAYES' THEOREM

So, coming back to our set of triplets, Eq. (1), we're expressing our state of knowledge about the frequency and damage, using probability curves. Probability curves are the language of uncertainty. Since the truth is, we always have uncertainty, we say that speaking in probability curves is telling the truth. Since knowing that truth is vital to the decision process, the risk analyst should always provide his input to the decision in the form of probability curves.

The next question is: Where do we get these curves? The answer is shown in Fig. 9; we get them from the evidence. Or you could better say, from the absence of evidence. If there was lots of evidence the curves would become spikes. These curves are determined by, dictated by, the evidence. But how are they dictated by the evidence? Answer; by Bayes' theorem. We list each item of relevant evidence and then process them, one by one, through Bayes' theorem. We call this the ''evidence-based'' approach. Initially, the curve is very fat, but as we add evidence items it homes in on the right answer.

OK, so what is Bayes' theorem? Bayes' theorem is the fundamental law of logical inference, i.e., the fundamental principle governing the process of evaluating evidence. This is what it looks like (see Fig. 10). "A" represents some hypothesis or proposition we're interested in. ''E'' represents the evidence we have relevant to this proposition. On the left side is our posterior probability, $p(A|E)$, i.e., our degree of confidence that A is true after we learn evidence E. On the right is $p(A)$, our probability prior to learning E.

6.1. Bayes' Theorem as the Definition of Logic

So the theorem tells us how much our confidence changes when we learn a new piece of evidence. This theorem has also been very controversial. What's the controversy? Many statisticians take the position that this theorem is formally true, like any other theorem in probability theory, but it's not much good for anything. To a Bayesian, this is not just another theorem. It's the fundamental law governing the evaluation of evidence.

To an extremist Bayesian, like myself, it goes deeper than that. It's not only the fundamental principle of logical inference, it's the very definition of logic itself. It's what we mean by logical, rational thinking.

In support of this extreme assertion, I refer you to Fig. 11, which shows that the fundamental principles of logic, the so called modus ponens and modus tolens, are

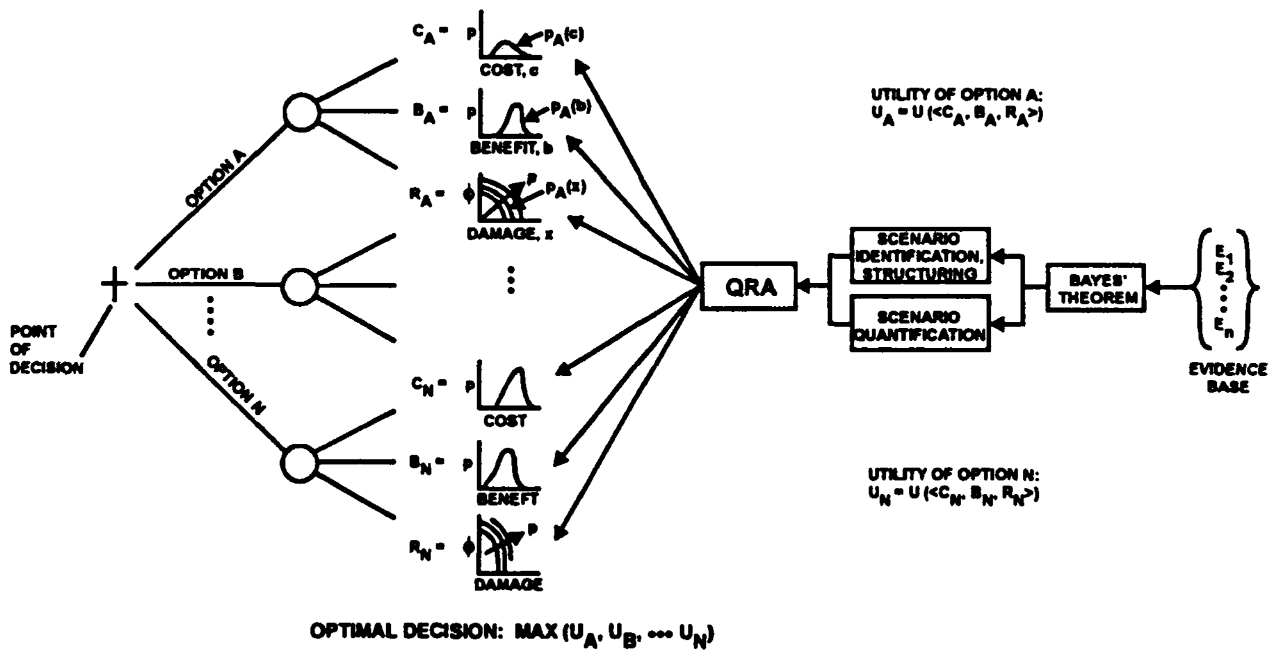

Fig. 12. The anatomy of a decision—the role of QRA and Bayes' theorem.

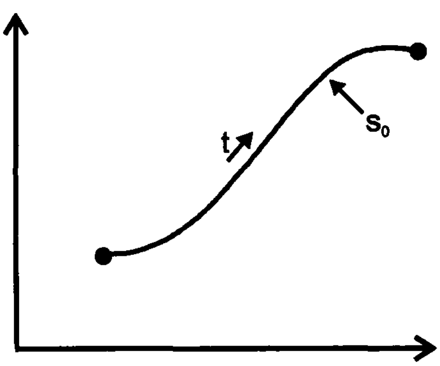

Fig. 13. Scenario s_0 viewed as a trajectory in the state space of the system.

special cases of Bayes' theorem. To me, this justifies adopting Bayes' theorem as the definition of logical thinking.

I find often that people react with surprise to the idea that the word "logical" needs to be defined. They seem to take the attitude that the supreme court justice took about pornography, saying "I can't define it, but I know it when I see it."

Note that this definition does not contain any value judgment. It's just a definition. If a particular piece of thinking conforms to this rule we call it logical. Otherwise we call it illogical, irrational, or incoherent. But that doesn't mean that it's necessarily "bad" or stupid. It just means that it doesn't conform. It doesn't say that logical thinking is necessarily better than any other kind. Indeed, we all know that sometimes it isn't.

7. OBJECTIVE/SUBJECTIVE PROBABILITY

Now with that introduction to Bayes' theorem, let's come back to the issue of defining probability (Table I). A Bayesian says probability means "confidence." But confidence exists only in our minds and is, therefore, subjective. The statistical school of thought objects to this, saying "that's unscientific. We don't want subjectivity, we want an objective science." So they rejected the whole Bayesian way of thinking, scornfully. The Bayesians answered with attacks of their own, and the battle has swung back and forth for two centuries.

But it's a misunderstanding. The whole two hundred years has been a miscommunication, caused by fuzziness in the words. Words like "confidence," and "belief" have a personal dimension to them. Are you a confident person; are you a believer, an optimist, a pessimist? Further down the list the words don't have that dimension. "Plausibility" and "credibility" are properties of the evidence, not of the person. A true Bayesian uses probability in that sense. Probability is that degree of credibility or confidence dictated by the evidence,

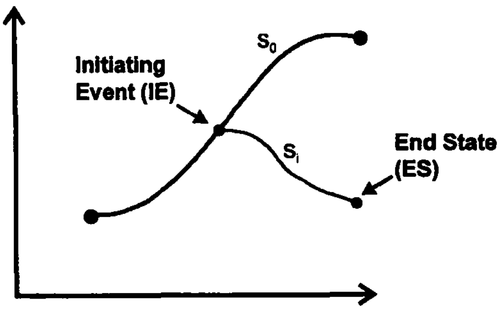

Fig. 14. The risk scenario s_i as a departure from s_0.

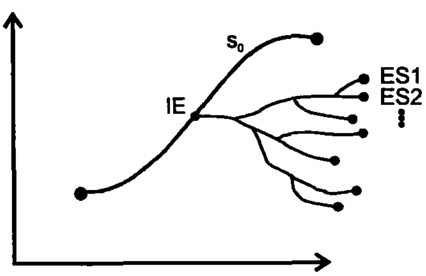

Fig. 15. Scenario tree emerging from the initiating event.

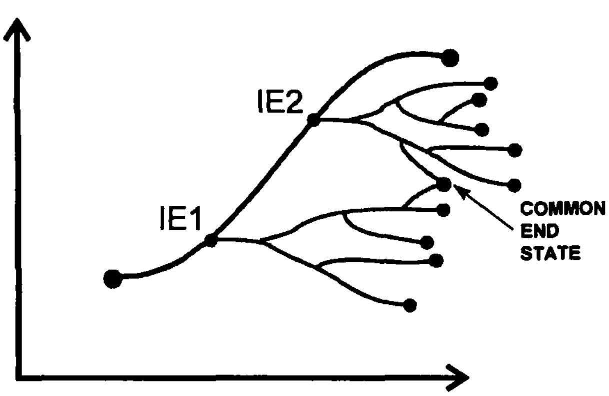

Fig. 16. Branches from two different trees can end at the same end state.

through Bayes' theorem. There's no personality in it, no "opinion."

The neatest statement of this point of view was given by Ed Jaynes as follows[7]:

> Probability theory is an extension of logic, which describes the inductive reasoning of an idealized being who represents degrees of plausibility by real numbers. The numerical value of any probability (A/B) will in general depend not only on A and B, but also on the entire background of other propositions that this being is taking into account. A probability assignment is "subjective" in the sense that it describes a state of knowledge

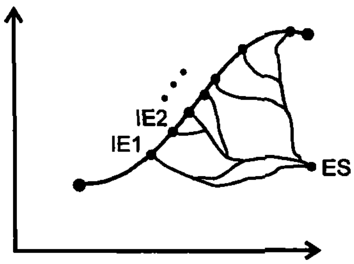

Fig. 17. "Incoming" scenario tree.

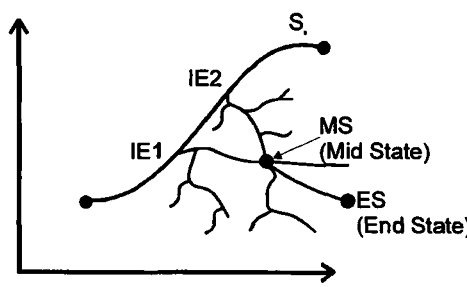

Fig. 18. In/out tree

> rather than any property of the "real" world; but is completely "objective" in the sense that it is independent of the personality of the user; two beings faced with the same total background of knowledge must assign the same probabilities.
>
> —E. T. Jaynes

The key point here is that while probability is "subjective" in that it measures something internal, namely degree of confidence, it can be defined to be entirely "objective," so that degree is determined totally by the evidence, and not by the personality or mood. And the way it is determined is through Bayes' theorem. So one could call this "objective/subjective" probability, if you like, or I like best the term "evidence-based" probability.

8. EVIDENCE-BASED DECISION MAKING

This idea of "objectifying" the so called subjective probability, has major implications. It resolves the historical controversies, and it shows us the how to put Risk Analysis on a totally solid conceptual foundation. It opens the way to what we can call "evidence-based" risk assessment and "evidence-based" decision-making. In regulatory and public decision making it shows us how, quantitatively, to

"Let the Evidence Speak!"

not the opinions, personalities, moods, politics, positions, special interests, or wishful thinking!

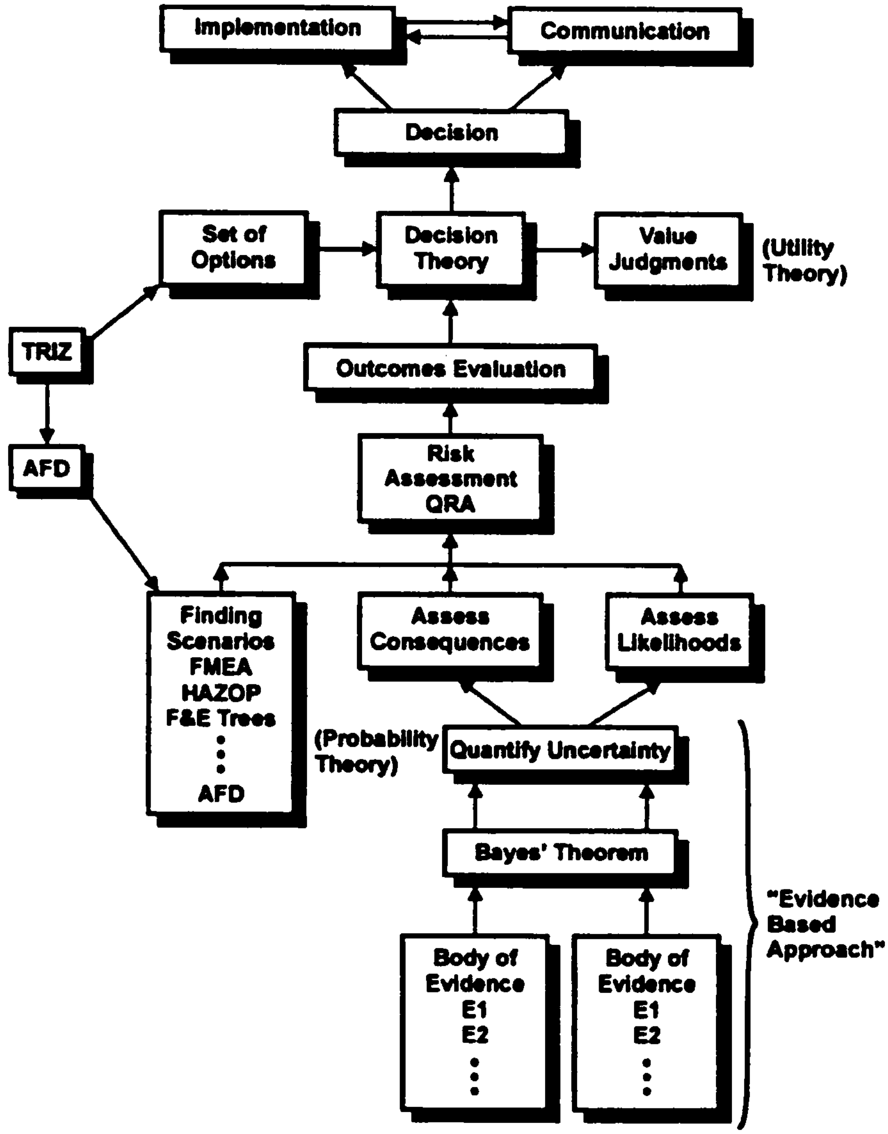

Fig. 19. How it all fits together.

By the way, this point of view guides us on how to deal with experts.[8] It tells us we should never ask an expert for his opinion. What we want from an expert is, his experience, his information, his evidence.

It also guides us on how to deal with a group of experts. We must first convert the question to a quantitative form, e.g., "What is the numerical value of this parameter F?" We then ask the experts "What evidence do we have relevant to this question?" As they answer we write down, and enumerate, a list of all the relevant evidence items available. We write down exactly what happened as distinct from our interpretations of what happened. We work over this list with the group until we obtain what we can call the "consensus body of evidence." At this point we apply Bayes' theorem to the list, item by item, at the end of which we have a consensus probability curve expressing what we collectively know about the value of F.

8.1. Structure of a Decision

Let's look at the structure of a decision problem (Fig. 12). At the point of decision we have to choose one of a set of available options. Each option brings with it costs, benefits and risks. Our knowledge of these at the point of decision is quantified in the form of probability curves. How do we calculate these curves? That is what QRA (quantitative risk assessment) is for. How does QRA do it? It does it by taking the whole body of evidence, the list of evidence items, and processing them through Bayes' theorem. So when the curves are obtained this way we call it "evidence-based" QRA, and when the decisions are made based on those curves, we call it "evidence-based" decision making.

8.2. Regulation

A decision problem also involves value judgments, represented in Fig. 12 by the utility functions. When it's a public decision, different parties will naturally have different utility functions, depending on how their interests are affected. So we, the public, employ regulatory agents to represent us in making sure that the decisions that get made reflect the interests of the community as a whole.

Regulators have a tough job. But they make it even more difficult, by trying to regulate using a "speed limit" concept. "Thou shalt have no more than so many ppms, or curies, etc." They then suffer and agonize over where to set this speed limit. They try to set it at such a level that the decision maker is forced to make the best decision for the community as a whole. Difficult to do. Worse, they attempt to set the limit without explicitly, and quantitatively, doing the decision analysis. That's next to impossible. No wonder they suffer. We are asking the wrong question. The question is not, "How much risk is acceptable?" The question is, "What is the best decision option?"

We need to get more of this decision theoretic point of view into regulation.

9. FINDING THE SCENARIOS

Returning to Eq. (1), we've talked about the last two parts of the triplet, but we haven't said anything about how to find the scenarios. Finding scenarios is part science and a large part art. Some useful ideas can be given, however. One is to think of S_0 as a trajectory in the state space of the system (Fig. 13). Any risk scenario, S_i, must then be viewed as a departure from S_0 (Fig. 14). There must therefore be a point of departure at which some "Initiating Event (*IE*)" happens. That starts the scenario and it goes until it ends at an "End State. (ES)"

Now actually (Fig. 15), from each initiating event, a whole treeful of scenarios emerges, depending on what happens next. This is called a "scenario tree."

Depending on how we define things, branches from two different trees can end up at the same end state (Fig. 16). That suggests that we could draw trees coming in to the end states of interest as in Fig. 17. This could be called an "incoming" scenario tree, also known as a "fault tree." The outgoing trees are also known as "event trees."

So this suggests two ways of finding scenarios. Method 1 is: find the IEs and draw the outgoing tree from each. Method 2 is: identify the end states of interest and draw the incoming trees to each.

There are also other methods; for example, identifying "middle states" from which we then draw both incoming and outgoing trees (Fig. 18).

9.1. Connection with TRIZ

In connection with finding scenarios, I'd like to introduce you to another new word. It's TRIZ, and it's an acronym for the Russian words meaning "Theory of the Solution of Inventive Problems." The Russians have been working on this theory for 50 years. It's a rather well developed thing.[9] What's of interest to us here is that there is a subsection of this theory, called Anticipatory Failure Determination (AFD), that applies to the problem of finding the scenarios.

It's interesting how they do it. Where QRA asks the question "What can go wrong?" The AFD people give this question an interesting twist. They ask, "If I wanted to make something go wrong, how could I do it?" In rephrasing the question this way they turn it into an inventive problem, and thus make available the whole apparatus of their theory, which is quite impressive. So I think that TRIZ has something significant to offer Risk Analysis.

10. PULLING IT ALL TOGETHER

Finally, I'd like to pull the pieces together in Fig. 19. The reason we do risk assessment is we have decisions to make. According to Decision Theory, to make a decision we need three things, a set of options from which to choose, an evaluation of the outcomes of each option, and a value judgment on each outcome. The role of QRA is to calculate those outcomes. And since we will always have uncertainty in the outcomes, we should, to tell the truth, quantify that uncertainty in the form of probability curves. For these curves to be worthy of trust, i.e., useful in a decision analysis, they should be based on the entire body of evidence available, evaluated through Bayes' theorem.

Now, Decision Theory is concerned with selecting the best option out of a given set. It says nothing about how we get those options in the first place. This is the creative part of the problem. Here tools like TRIZ can be helpful to invent new and better options. They can also be helpful, as we said before, in "inventing" scenarios that we may not have thought of otherwise.

Finally, making the decision is not the end of the job. It's necessary to get the decision accepted and implemented. For that we need the support of the people affected by it. That means risk communication, and decision communication. For that to take place, it's crucial that we have words that we all understand and use in the same way. That is what this talk has been about. So thank you again for the opportunity to speak to you, and I hope it's been useful.

REFERENCES

1. C. Howson and P. Urbach, "Scientific Reasoning: The Bayesian Approach, " Open Court, LaSalle, Illinois, 1989.
2. B. de Finetti, "La prévision; ses lois logiques, ses sources subjectives," *Ann. Institut Henri Poincaré* **7**, 1–68 (1937) (reprinted in English translation as "Foresight: Its Logical Laws, Its Subjective Sources," in H. E. Kyburg, Jr. and H. E. Smokler (eds.), *Studies in Subjective Probability* (Wiley, New York, 1964); *Theory of Probability* (Wiley, New York, 1979).
3. H. Jeffreys, *Theory of Probability*, 3rd Ed. (Clarendon Press, Oxford, 1961).
4. K. J. Arrow, *Essays in the Theory of Risk-Bearing* (American Elsevier Publishing Company, New York, 1976).
5. B. Kosko, "Fuzziness vs. Probability," *Int. J. Gen. Syst.* **17**(2), 211–240 (1990).
6. S. Kaplan and B. J. Garrick, "On the Quantitative Definition of Risk," PLG-P0196, *Risk Anal.* **1**(1) (1981).
7. E. Jaynes, Notes to a short course at University of California–Los Angeles (July 1960).
8. S. Kaplan, " 'Expert Information' versus 'Expert Opinions'; Another Approach to the Problem of Eliciting/Combining/Using Expert Opinion in PRA," PLG-P0682, *J. Reliab. Eng. Syst. Safety* **35**, 61–72, (1992).
9. S. Kaplan, *An Introduction to TRIZ, The Russian Theory of Inventive Problem Solving* (Ideation International, Southfield, MI, 1996).