Vol. 43, No. 6 • Nov.-Dec. 2007

What Do You Really Know?

by Ira J. Rimson and Ludwig Benner, Jr.

Page: 1 | 2 | 3 | 4

The greatest obstacle to discovering the shape of the Earth, the continents and the oceans was not ignorance, but the Illusion of Knowledge.
— Daniel Boorstin, Librarian of the U.S. Congress, Emeritus

In his paper "What is Missing? A Philosophy of Risk Management,"[1] Michael Murphy[2] argues that high-cost risk management failures might be mitigated, were risk managers to give greater credence to the major traditional branches of philosophy. He listed 15 well-known cases of risk management failures that occurred in recent history, from the Chevrolet Corvair to the Walkerton, Ontario, e. coli outbreak, to Enron, and enumerated the influences of lack of consideration for philosophic precepts in their systems' ultimate collapse. Among the five major philosophical branches, two are represented in all cases on Murphy's list: Metaphysics and Logic. Metaphysics is comprised of Ontology (what exists in reality vs. what does not; *i.e.*, reality vs. perception) and Epistemology (the theory of knowledge, its limits and the validity of what we know). Logic defines the validity of inference and reason, which overlies the entirety of system development. We have chosen to concentrate this column on the issue of epistemology because, in our experience, it is the philosophical discipline most commonly violated by risk managers and analysts. It is also likely the easiest to interject into current system safety practices.

Recent risk-management failures have continued to demonstrate the need to expand traditional analytic methodologies to include "what did we know?"; "what did we need to know?"; and "how accurate and valid were our data and the assumptions we derived from them?" We have emphasized in prior columns that humans are the principal — and least predictable — variables in systems' operation. The risks inherent in inanimate systems are calculable. The infinite intricacies of human behaviors bias and change those risks beyond capability for reliable mathematical calculation. Such simplistic explanations as "It was an error," or "We didn't anticipate that [X] would respond that way," are useless for establishing any kind of behavioral change that might preclude recurrence.

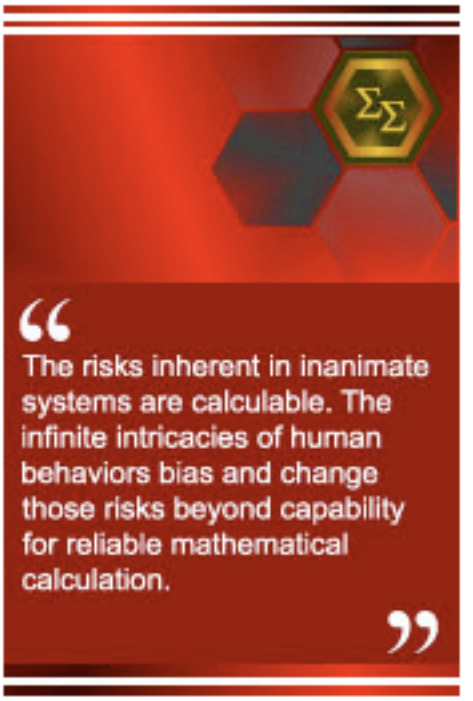

Epistemology is the philosophy of knowledge, and is useful for determining what we *need* to know, how much we really know compared to what we need to know,[3] and the degree of *uncertainty* of what we think we know.

The London School of Economics' *Risk and Regulation Magazine* recently published a paper by Nicholas Taleb and Avital Pilpel titled "Epistemology and Risk Management," in which they argue:

> *…[T]he production of a risk "measure" must be subjected to the question "how do you know what you claim to know" .... Claims regarding risk cannot be made without any rigorously established supervision of their validity. There is a need for skeptical inquiries concerning how a risk measure was obtained and how an opinion was formed.*[4]

We believe that absence of skeptical inquiry is a critical inadequacy of current system safety practice. Skepticism is particularly significant when risk analyses are derived from historic data that analysts assume to be accurate predictors of future events that might bear little, if any, resemblance to the assumed risk inputs. Analysts don't often delve into the details of their assumed "precursors," least of all into behavior of human players in the system. Most system safety analyses relegate human operator behavior to a binary mode: either the operator(s) perform according to preprogrammed expectation — doing what they are supposed to — or they don't — *not* doing what they are supposed to. Analysts rarely consider the infinite variety of alternative responses that human behavior *can* instantiate in the system.[5]

next page »

Worse yet, when unpredicted undesired events eventually occur (and they usually do), few analysts return to their original inputs and premises to find out how and where their predictions failed. They compound their original inaccuracies downstream to the next naïve analyst who assumes, without verification, the implied accuracy of the original prediction(s). It often seems that once an analysis has been calculated, any motivation to revise it to conform to reality disappears. It's more likely that errant analysts will dig out some past happening and point to the latest events as "retrocursors," by which to establish the prophecy of their original analyses.

In the referenced paper, Taleb and Pilpel summarize their First Epistemological Problem: "induction and small probability," thus:

> *If small probability events carry large impacts, and (at the same time) these small probability events are more difficult to compute from past data itself, then: our empirical knowledge about the potential contribution — or role — of rare events (probability x consequence) is inversely proportional to their impact. … We understand so little about catastrophic events, yet these are the events that we talk about the most casually. In risk management terms, the bigger the event, the less we have a clue.*

Risk management should attempt to identify and minimize the effects of high-cost events, irrespective of their probability of occurrence. High-probability events are relatively easy to spot, so most efforts are concentrated on low-probability events. In practice, risk analysts "…confidently extrapolate from the seen to the unseen…"[6] without bothering to establish their equivalence. Perhaps system safety practitioners should bother themselves with the equivalence issues of the inputs.

Inputs to risk analysis include both actual historical occurrence data and estimated data. Actual data are derived from documentation of past events. Risk is often analyzed with logic trees showing the progression of coupled events and conditions as inputs leading to some outcome, with numerical probability values assigned to the coupled elements and outcomes based on actual occurrence data, or "expert" estimates based on estimated data. Our challenge: Are those risk inputs and resultant probability outputs accepted uncritically? Do analysts thoroughly examine the epistemology of the problem?

Paraphrasing Yogi Berra,

> *Good past data has a lot of good in it, but it's the bad side that's bad.*[7]

How do analysts determine whether what they think they know — the data they use — belongs to the "good side" or the "bad side?" Were analysts to substitute the probability <P=1> for fractional probabilities in all logic tree analyses, would the outcome require further analysis to determine the real risk?[8]

More attention needs to be focused on answers to the questions "what do we *need* to know"? and "what is the *validity* of what we think we know?" Boorstin's "Illusion of Knowledge" is the product of assumptions multiplied by ignorance. One need merely compare the plethora of contradictory "scientific" studies on the effects of foods, medications, exercise, stress, caffeine, global warming and/or cooling, …*ad infinitum*, to realize the prevalence of illusory knowledge in our society.[9] The "good" thing is that the probability that these common manifestations of uncommon ignorance might result in fatal or highly costly outcomes is no better than chance. The "bad" is that the probability of their occurrence is >P=0, and no analyst can predict accurately when that shoe will drop.

The author of a letter published in the most recent e*JSS*[10] disagreed with our column in V. 43, #2, in which we described the COMAIR flight crew's take-off decision at Lexington, Kentucky, as an error in judging the criticality of the potential outcome(s) of their behavior choice. His arguments typify the absence of epistemological examination common in risk analyses:

> *I do disagree with a few of the details, such as their assertion that the crew of the airliner on the wrong runway "didn't recognize the potential criticality of the outcome." Of course the crew knew the potential criticality of trying to take off on a short runway, they just didn't know they were doing such a thing. They didn't know that the runway was far too short and decided to give it a go anyway. They were mistaken — they made an error, but not an error in judging the criticality of their decision.*

What does any analyst *need to know* in order to reach a valid conclusion regarding human behavior in given circumstances? At the outset, the analyst should know what happened. In addition, it is highly advisable to verify that the facts upon which any assumptions will be based are valid. Next, it helps to know what you *really* know. The author asserts in this case not only that he knew what the crew *knew*, he also assumes what the crew *did not* know, ignoring the impossibility of proving a negative.[11]

The writer's argument is tautology — a circuitous repetition of unverified assumptions, each feeding from its predecessor, as to the crew's knowledge (or lack of same).

If we accept the writer's assumption that "…the crew knew the potential criticality of trying to take off on a short runway," then they should never have attempted the take-off without accurately ascertaining their position on the airfield. Logic suggests that, if they *knew* the criticality of the wrong decision, then they would have ensured that they were on the correct runway.[12]

On the other hand, the writer asserts that the crew "…just didn't know they were doing such a thing." If the crew *knew* the potential criticality of attempting to take off on the shorter runway, but chose not to ascertain that they were on the longer one and "gave it a go anyway," then the crew not only violated FAA regulations and their training, they *deliberately* risked 50 lives and fatally misjudged the criticality of their decision — if they considered it at all.

Whatever the motivation for their behavior, and whatever you choose to call it, we stand by our contention that the crew made a fatal error in judgment. It resulted in 49 unnecessary deaths, and possibly the demise of the airline.

We invite arguments pro or con from our readers.[13]

[11] The author could not have studied the NTSB's "Probable Cause" statement because it was not released publicly until late July, 2007, well after the issue went to press.

[12] This is especially significant in this case, as there are two runways that originate in the same vicinity on the airfield: one is sufficiently long to have accommodated the attempted takeoff; the other is not.

[13] Thanks to Mike Murphy, who provided epistemological review and input.