

Risk Management in Complex Organizations

Richard Bookstaber

We can manage risks only when we can identify them and ponder their possible outcomes. We can manage market risk because we know security prices are uncertain; credit risk because we know companies can default; operational risk because we know missteps are possible in settlement and clearing. But despite all the risks we can control, the greatest risks remain beyond our control. These are the risks we do not see, things beyond the veil. The challenge in risk management is how to deal with these unidentified risks. It is more than a challenge; it is a paradox: How can we manage a risk we do not know exists? The answer is that we cannot manage these risks directly, but we can identify characteristics of risk management that will increase our ability to react to the risks. These characteristics are easiest to grasp in the biological setting, where we are willing to concede that nature has surprises that are wholly unanticipated by life forms that lack total foresight. A disease that destroys a once-abundant food source, the introduction of chemicals into a pristine environment, the eruption of a volcano in a formerly stable geological setting are examples of events that could not be anticipated, even in probabilistic terms, and therefore, could not be explicitly considered in rules of behavior. They are nature's equivalent to the unforeseeable risks for the corporation.

The best measure of adaptation to unanticipated risks in the biological setting is the length of time a species has survived. A species that has survived for hundreds of millions of years can be considered, *de facto*, to have a better strategy for dealing with unanticipated risks than one that has survived for a short time. In contrast, a species that is prolific and successful during a short time period but then dies out after an unanticipated event may be thought of as having a good mechanism for coping with the known risks of one environment but not for dealing with unforeseeable changes.

By this measure, the lowly cockroach is a prime case through which to study risk management.

Because the cockroach has survived through many unforeseeable changes—jungles turning to deserts, flatland giving way to urban habitat, predators of all types coming and going over the course of the countless millennia—the cockroach can provide us with a clue for how to approach unanticipated risks in our world of financial markets. What is remarkable about the cockroach is not simply that it has survived for so long but that it has done so with a singularly simple and seemingly suboptimal mechanism: It moves in the opposite direction of gusts of wind that might signal an approaching predator. This “risk-management structure” is extremely coarse; it ignores a wide set of information about the environment—visual and olfactory cues, for example—which one would think an optimal risk-management system would take into account.

This same pattern of behavior—using coarse decision rules that ignore valuable information—appears in other species with good track records of survivability. For example, the great tit does not forage solely on the small set of plants that maximize its nutritional intake. The salamander does not fully differentiate between small and large flies in its diet. This behavior, although not totally responsive to the current environment, allows survivability if the nature of the food source unexpectedly changes. We also see animals increase the coarseness of their response when the environment does change in unforeseeable ways. For example, animals placed for the first time in a laboratory setting often show a less fine-tuned response to stimuli and follow a less discriminating diet than they do in the wild.

The coarse response, although suboptimal for any one environment, is more than satisfactory for a wide range of unforeseeable environments. In contrast, an animal that has found a well-defined and unvarying niche may follow a specialized rule that depends critically on that animal's perception of its world. If the world continues on as the animal perceives it—with the same predators, food sources, and landscape—the animal will survive. If the world changes in ways beyond the animal's experience, however, the animal will die off. Precision and focus in addressing the known comes at the cost of reduced ability to address the unknown.¹

Richard Bookstaber is head of risk management at Moore Capital Management.

What does this concept mean for risk management in the corporation? Imagine that the CEO of a corporation has been given an omniscient view of the world, so he can see all possible sources of risk. Armed with this knowledge, he is then allowed to design an approach to risk management but with one constraint: Although aware of all the risks that the company has not identified, the CEO is not permitted to convey any information about these risks to others in the company. That is, he can create a structure to address both the seen and the unseen risks, but he cannot open up the unseen to view. He can use full information to create the structure, but the resulting structure cannot convey information that would not otherwise be known. How will the risk-management structure he designs differ from the structure he would design if he were never afforded this omniscient view?

If he follows the lessons from the biological world, he will move resources away from dealing with details of the known risks and reconfigure the risk-management structure for better response to the risks that remain unknown. Knowing that the structure cannot address all the risks, he will streamline and simplify formal processes and procedures because some of those procedures will obscure the unseen risks. He will decide that it is better to spend less time focusing on detailed investigation of the known risks and more time thinking and reacting to the unknown risks. Similarly, he will simplify the risk-management models and analyses. Specialized analysis, although important in providing perspective for what is known, can only coincidentally do the same for the unknown. Coarse measures will be more likely to indicate—although perhaps not fully elucidate—areas of unanticipated risk. And by being more concise, these measures will be easier to discuss and analyze intuitively than specialized measures. He will reduce the organizational complexity and hierarchy of responsibility. This complexity, although perhaps effective in a specific environment, will obscure unidentified risks that fall across organizational lines and slow the company's ability to respond to events that are beyond the intended design and function of the risk-management system. In short, this CEO will eschew pinpoint laser targeting of the observed risks in favor of lower-resolution, 360-degree radar that is more likely to capture the unobserved risks.

The unanticipated events that are catastrophic for a species are not difficult to understand after they have happened. We can see that a species was simply not designed to anticipate and react to certain events. As recent history has shown, the same is true of catastrophic risks in the financial world.

Whether it is Daiwa Bank, Sumitomo Bank, Kidder Peabody & Company, or Barings Securities, now, the catastrophic risk can easily be described in one or two sentences (and described without resorting to Greek letters or statistics). That the risks are obvious after the fact is not surprising; after all, every loss can be explained *ex post*. It is tempting to think "they should have known" or "how could anyone run a firm like that!" But simple as the events turn out to be—and simple as the changes in management or oversight needed to eliminate the problems—new events continue to catch us unaware.

The root of the problem is not the complexity of the unseen risks. The unforeseen events turn out to be simple. The root of the problem is the complexity of the organization. This relationship between organizational complexity and risk is well understood in two industries in which a risk-management failure is measured in lives rather than in dollars: airlines and nuclear energy.

The airline industry has three categories of accident, and each of these categories has a counterpart in investment management and trading organizations.² The first category, procedural accident, covers accidents arising from mistakes in following obvious procedures, such as piloting without adequate sleep or flying without the requisite reserves of fuel. On the trading floor, procedural errors tend to be operational ones, such as an incorrect input or miscommunication of an order. Among the standard procedures for addressing these sorts of risks are coding buy and sell orders in different colors and communicating orders in full detail.

The second category, engineering accident, comprises accidents stemming from structural problems that, in principle, should have been discovered during flight tests or routine checks but were not. On the trading floor, engineering-type accidents lie largely in the realm of market risk and range from model misspecification to incomplete representation of risk factors.

The last category, system accident, covers accidents that could not be anticipated and addressed through detailed procedures or engineering. They are born from the confusion and interrelationships inherent in complex organizations and systems. They have been termed "normal accidents" because they are the inevitable, unavoidable "toxic waste" that comes with building and running complex systems (Perrow 1984).

The dilemma of normal accidents is this: Operate with a given level of complexity, and over time, you will observe a given number of normal accidents; add safeguards to try to reduce these acci-

dents, and you add complexity that will increase them. Indeed, the safeguard itself may become the source of the accident. The ValuJet crash occurred, for example, because oxygen generators, a required safety back-up, were being removed to comply with yet other safety requirements. The near-disaster at Three Mile Island occurred because the safety system was so complex it became oppressive and was, accordingly, ignored. The Chernobyl disaster occurred because of runaway events triggered by a mandated safety test.

The recent mergers in the financial industry have moved system complexity several steps up the scale. The lesson of the cockroach becomes that much more germane. The finely tuned approach to dealing with risk—the approach that would seem optimal in any one world—may in the long run prove suboptimal. Given the complexity and fundamental unpredictability of nature, an approach that is coarse and less complex may be the best long-term risk-management strategy. The lesson for the corporation or investment firm is even more pointed: Unlike the biological world, in the business world, the more complex risk-management rule may actually make the situation worse, leading to greater complexity rather than simply a less-than-ideal response.

Once a normal accident occurs, it comes through the veil and is open to view. It is then manageable, and controls can address recurrences. But if a system is already at a level of complexity where normal accidents are common, adding one control after another will simply increase complex-

ity and obscurity. Because increased complexity increases the likelihood of normal accidents, this response will be counterproductive. Indeed, the conventional response to the unforeseeable events will produce a system “risk cycle” that will lead to better and better control of the identified risks but an upward spiral of surprise events.

The dollar cost of the increase in controls will be immediately evident. The costs that will not be evident at first (indeed, may never be evident) are those born of increased complexity and the reduced “survivability” of the organization in the face of the unforeseeable risks. If these costs were recognized, a normal accident would not simply change the controls in the risk-management structure; it would change the risk-management structure itself. That is, the realization that risk management may fail in an unanticipated way provides information that is beyond the scope that new controls alone can address. The unforeseeable events provide information about the incomplete nature of any particular risk-management paradigm. The knowledge that risks beyond the scope of the model remain should lead to a reconsideration of the information used for monitoring risk and to a broader, not a more detailed, response.

This article is based on presentations given to the Berkeley Finance Symposium and the Federal Reserve Bank of Atlanta Risk Management Conference.

Notes

1. The analysis of optimal animal behavior under extended uncertainty is presented in Bookstaber and Langsam (1985). The article provides more examples of behavior that appears to be suboptimal given the animal's view of the world but that is superior to a fine-tuned rule when unfore-
2. The delineation of the categories and a general discussion of the limitations of assuring safety in a complex system such as the airline industry are presented in Langewiesche (1998).

References

- Bookstaber, Richard, and Joseph Langsam. 1985. “On the Optimality of Coarse Behavior Rules.” *Journal of Theoretical Biology*, vol. 116:161–193.
- Langewiesche, William. 1998. “The Lessons of ValuJet 592.” *Atlantic Monthly*, vol. 281, no. 3 (March):81–97.
- Perrow, Charles. 1984. *Normal Accidents*. New York: Basic Books.