

## Joshua Rosenberg: Why do risk events occur? Insights from accident models

Remarks by Mr Joshua Rosenberg, Executive Vice President and Chief Risk Officer of the Federal Reserve Bank of New York, at the 7th Annual Risk Americas 2018 Conference, New York City, 17 May 2018.

\* \* \*

Thank you for the opportunity to present one of today's keynote addresses. Before I continue, I should note that the views presented here are my own, and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

Two of the most powerful questions in risk management are "What if?" and "Why?" My remarks will focus on the second question, "Why?" and more specifically on the question "Why do risk events occur?"

Risk events -failures that impair an organization's ability to achieve its objectives - provide a unique window into an organization's risk and control environment.<sup>1</sup> Because risk events represent a translation of the possibility of risk to actuality of risk, we have much to learn from their causes. By understanding causes, we can improve our ability to design and implement controls to prevent, detect, and correct problems before they mature into risk events. In addition, after a risk event occurs, understanding causes helps us extract the right lessons, not just to fix the immediate issue, but also to address any underlying vulnerabilities.

My intention in this presentation is to explore the causes of risk events through the lens of models of organizational accidents. Based on the work of Rosness et. al.,<sup>2</sup> Dekker et. al.,<sup>3</sup> and Yang and Haugan,<sup>4</sup> I select four accident models: the barrier model, the information processing model, the conflicting objectives model, and the normal accident model. In my commentary, I draw heavily on their interpretation and analysis, but of course, any errors or misinterpretations are my own. While the accident literature also has important lessons for us on risk detection and risk response, I'll focus my remarks on risk prevention. Along the way, I'll highlight what we can take from these models to improve our ability to understand and manage enterprise risks.

I start with the barrier model, which we can consider as a baseline model for understanding accidents.<sup>5</sup> In the barrier model, an accident occurs when a hazard penetrates protective barriers and causes harm. Barriers can be "hard," like a guardrail placed between people and dangerous equipment, or "soft," like safety rules, training, or management oversight. In James Reason's elaboration of the barrier model, the ideal condition of a barrier is solid, but the reality is that barriers have weaknesses or holes like a slice of Swiss cheese.<sup>6</sup> These weaknesses can arise from human, technical, and organizational factors. They can be the result of design flaws or problems with implementation.

What does the barrier model imply for risk event prevention? The barrier model tells us that risk events occur when threats penetrate barriers, so to prevent risk events we should either reduce threats or strengthen the barriers.<sup>7</sup> This is familiar territory for risk managers, although we usually use the term "control" rather than "barrier." Strengthening barriers in this context means designing and implementing effective controls, and then maintaining and monitoring those controls to make sure they continue to operate as intended. We are also quite familiar with the approach of establishing multiple barriers (or layered defenses) to provide more protection than a single barrier.

The second model is the information processing model developed by Barry Turner.<sup>8</sup> It provides a different perspective on the causes of organizational accidents and, I would argue, important

insights for enterprise risk management. While the barrier model evokes a physical connection between hazards and accidents, the information processing model is set in the context of an organization's own understanding - or misunderstanding - of its operations and risks. A key insight of the model is that what is usually considered a strength of large organizations - their ability to create a shared outlook and direct common effort - becomes a weakness when the shared outlook results in collective blindness to problems.

In this model, an accident occurs when an organization is unable to perceive, and therefore does not respond to, a long string of early warning signals that point to increasing danger. Here, accidents are "man-made disasters" that are the result of a failure of organizational intelligence. A frequently-cited application of this model is Diane Vaughan's analysis of the Challenger Space Shuttle disaster.<sup>9</sup> In that work, she coins the term "normalization of deviance" to describe how organizations can grow accustomed and unresponsive to signs of problems. One of my colleagues refers to this as noseblindness: that is, you will eventually become accustomed to any smell, no matter how bad it is.

What does this imply for preventing risk events? In the information processing model, accidents are caused by failures of knowledge, perception, and communication, so that is also where the solutions lie. To avoid these risks, we should maintain a deep, broad, and current risk picture that is sensitive to changes in the internal and external environment. We should encourage a culture of vigilance and challenge of accepted beliefs, including seeking alternative, independent, and outside perspectives. And, our organizations should ensure that there are effective channels for timely risk communication and escalation.

The third model, Jens Rasmussen's conflicting objectives model, also identifies the organization itself as a source of risk, although in this case the focus is on incentive conflicts rather than information gaps.<sup>10</sup> Accidents here are the end result of the incremental tradeoffs that managers and staff make every day as they balance competing demands on their time and attention. Because organizations face pressure to increase efficiency, many individuals make choices in their own work to favor production over safety.

While each of these choices in isolation is innocuous, the incremental effect of each choice pushes the organization, step by step, beyond the boundary of safe operations and into an uncontrolled state that leads to an accident. Rasmussen suggests that several large-scale organizational accidents, including the Chernobyl nuclear disaster, were caused by "systematic migration of organizational behavior toward accident under the influence of cost-effectiveness in an aggressive, competitive environment."<sup>11</sup> It is not only cost that drives conflicting interests; career management and behavioral biases can also lead individuals to make decisions that favor short-term outcomes over the long-term interests of the organization.

What are the lessons here for the prevention of risk events? In the language of this model, an organization should ensure that safety and performance priorities are balanced, that it is aware of where the boundary of safe operations is, and that it responds when it is too close to the boundary. This advice is in line with guidance and practice on the use of risk appetite as a core component of effective enterprise risk management.<sup>12</sup> By articulating risk appetite and monitoring risk against appetite, organizations are better equipped to see their risks and manage them within appropriate boundaries.

There is an interesting connection between the last two models, both of which emphasize cultural and behavioral risks, and the concept of risk culture. In my view, a sound risk culture, as defined for example in the Financial Stability Board's (FSB) risk culture framework, can address many of the risk drivers identified by these accident models.<sup>13</sup>

The FSB framework states that a "sound risk culture promotes an environment of open communication and effective challenge in which decision-making processes encourage a range

of views; allow for testing of current practices; [and] stimulate a positive, critical attitude among employees.”<sup>14</sup> This should help organizations avoid failures of information processing. Furthermore, the FSB explains that in a sound risk culture, “[p]erformance and talent management encourage and reinforce maintenance of the financial institution’s desired risk management behaviour.”<sup>15</sup> So, by establishing incentives to emphasize effective risk management, an organization can lean against the pressures highlighted in the conflicting objectives model.

The normal accident model, developed by Charles Perrow, is the fourth model that I would like to share with you.<sup>16</sup> This model focuses on complexity, interdependence, and opacity in the interplay between people, processes, and systems in large organizations. Here, large-scale accidents are a rare but normal outcome of multiple failures that interact in unpredicted ways. Perrow analyzes the nuclear meltdown at the Three Mile Island power plant as a case study of a normal accident.

In the normal accident model, barriers help prevent failures of individual components, but when barriers add to overall system complexity and opacity, they can contribute to major system failures. Ironically, the defenses that we establish to protect ourselves against small problems become the threats that can cause even larger problems. An important insight from the normal accident model is that systems with two characteristics—high interactive complexity (likely to exhibit unexpected sequences of events) and tight coupling (where a change in one component leads to a rapid and strong change in other components) - are at greatest risk of systemic failure. Systems with both of these characteristics, in Perrow’s view, are destined to experience catastrophic accidents, in spite of our best efforts to control them.

What do we learn when we look at risk in our organizations from the perspective of this model? One message is that we should simplify and streamline processes, especially those that are interactively complex and tightly coupled. We can look for and address instances where layered controls add, rather than reduce, risk due to an increase in overall complexity. More generally, organic accumulations of controls should be replaced with intentionally designed control systems. For example, frameworks, like Carnegie Mellon’s Resilience Management Model, can be used to harmonize risk management activities across the enterprise.<sup>17</sup>

After reviewing these four models, a sensible question to ask is, “Which one is right?” As a response, I will give you an economist’s answer: “it depends.” In my view, each of these models has lessons to teach, and which model is the most useful will depend on the context in which it is applied.<sup>18</sup> In my view, the barrier model provides a practical approach to identify and analyze risks and controls. By broadly interpreting barriers to include aspects of organizational structure, culture, incentives, and information flow, I think the differences shrink between the barrier model, the information processing model, and the conflicting objectives model.<sup>19</sup>

As I said at the start, one of the reasons we want to understand why risk events occur is to improve our ability to prevent them. To go a bit further with this, let’s modify the question “Why do risk events occur?” and instead ask, “Why did this specific risk event occur?” That leads us to the topic of learning from risk events. Here, we are looking not only to learn how to prevent the same failure from happening again, but also to find and remediate related vulnerabilities across the organization that would lead to future failures.

When we investigate a risk event, we often look to root cause analysis as a tool to understand why. Root cause analysis fits well into the barrier model perspective; we have found the root cause of a risk event when we have identified the barriers that failed. And, importantly, the accident models give us a rich set of potential causes of a risk event to consider. The information processing and conflicting objectives models pull us away from the incident itself to look at broad organizational factors. The normal accident model leads us to question whether it is even possible to identify a clear cause amidst the tangled interactions of a complex system.

One thing that I have learned from the accident models is not to quickly or easily accept the answer that human error is the root cause of a risk event. James Reason makes a strong case that, while a human is likely to be the last and most visible link in the chain of an accident, human error is rarely the underlying cause. Reason draws our attention to work process and organizational causes such as “unworkable procedures, clumsy automation, shortfalls in training, less than adequate tools and equipment” and culture.<sup>20</sup>

In further defense of people, Hollnagel et al. introduce the concept of Safety I and Safety II.<sup>21</sup> Under Safety I, the traditional view of safety, people are treated as hazards whose unreliability leads to accidents. Instead, Hollnagel et al. advocate a perspective, called Safety II, in which people are seen as mitigants in that they contribute flexibility and resilience to processes.

How an organization learns from its mistakes - Is an independent investigation performed? Do lessons learned lead to change? - is itself a reflection of its risk culture. I suggest that, when we are investigating risk events, we widen our field of vision using the perspectives of the accident models. As a starting point, here are some questions that might be worth adding to your repertoire.

- ♦ To look beyond the failure of a single barrier: Why did the control system - including preventive, detective, and corrective controls - fail?
- ♦ To look for gaps in information processing: Were there early warning signs of the event, and, if so, how did the organization respond? Were there effective processes in place to detect and address organizational blindspots?
- ♦ To look for conflicting objectives: How did processes and controls evolve over time in response to incentives and resources?
- ♦ To look for interactive complexity and tight coupling (causes of normal accidents): Did multiple failures interact and amplify each other in unexpected ways? Was the complexity of the system itself part of the cause?

To close, I'll draw your attention back to two effective tools we already have in our toolkit to address risks highlighted by organizational accident models: managing risks within risk appetite and reinforcing a sound risk culture. We should take full advantage of their potential.

---

<sup>1</sup> See, for example, Committee of Sponsoring Organizations of the Treadway Commission, 2017. “Enterprise Risk Management: Aligning Risk with Strategy and Performance.”

<sup>2</sup> Rosness, R., Grøtan, T. O., Guttormsen, G., Herrera, I. A., Steiro, T., Størseth, F., Tinmannsvik, R. K., and Wærø, I., 2010, “Organisational Accidents and Resilient Organisations: Six Perspectives,” Revision 2, SINTEF Report A17034.

<sup>3</sup> Dekker, S. A., Hollnagel, E., Woods, D. D. and Cook, R., 2008. “Resilience Engineering: New Directions for Measuring and Maintaining Safety in Complex Systems,” Final Report, Lund University, School of Aviation.

<sup>4</sup> Yang, X.; Haugen, S., 2014. “A Fresh Look at Barriers from Alternative Perspectives on Risk,” Probabilistic Safety Assessment and Management,” pp. 2014–06-22 – 2014–06-27.

<sup>5</sup> Haddon, W., 1980. “The Basic Strategies for Reducing Damage from Hazards of All Kinds,” Hazard Prevention, Volume 16, Issue 1, pp. 8-12.

<sup>6</sup> Reason, J. 1997. Managing the Risks of Organizational Accidents, Brookfield: Ashgate Publishing Company.

<sup>7</sup> Going beyond strong barriers, resilience engineering emphasizes the importance of an organization's ability to adapt and maintain operations in the face of disruption. See, for example, Hollnagel, E., Woods, D.W., and Leveson, N., 2006. *Resilience Engineering: Concepts and Precepts*, Burlington: Ashgate Publishing Company.

<sup>8</sup> Turner, B.A., 1978. *Man-made Disasters*, London: Wykeham Science Press.

<sup>9</sup> Vaughan, D., 1996. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*, Chicago: University of Chicago Press.

<sup>10</sup> Rasmussen, J., 1997. "Risk Management in a Dynamic Society: A Modeling Problem," *Safety Science*, Volume 27, Issue 2-3, pp. 183–213.

<sup>11</sup> *Ibid*, p. 189.

<sup>12</sup> Rittenberg, L. and Martens, F. 2012. "Understanding and Communicating Risk Appetite." Committee of Sponsoring Organizations of the Treadway Commission.

<sup>13</sup> Financial Stability Board, 2014. "Guidance on Supervisory Interaction with Financial Institutions on Risk Culture: A Framework for Assessing Risk Culture." Also, see Institute of Risk Management, 2012. "Risk Culture: Resources for Practitioners."

<sup>14</sup> *Ibid*, p. 4.

<sup>15</sup> *Ibid*.

<sup>16</sup> Perrow, C. 1984. *Normal Accidents: Living With High-Risk Technologies*, New York: Basic Books.

<sup>17</sup> Caralli, R.A., Allen, J., White, D.W., Young, L.R., Mehravari, N., and Curtis, P.D., 2016. "CERT® Resilience Management Model, Version 1.2," CERT Program, Carnegie Mellon University.

<sup>18</sup> Applied accident investigation techniques are compared in Sklet, S., 2004. "Comparison of Some Selected Methods for Accident Investigation," *Journal of Hazardous Materials*, Volume 111, pp. 29–37.

<sup>19</sup> A systems approach to analyzing accidents that incorporates a broad range of causal factors is developed in Leveson, N., 2004. "A New Accident Model for Engineering Safer Systems. *Safety Science* Volume 4, Issue 2, pp. 237–270. The resilience engineering approach, cited in footnote 7, is another way to synthesize insights across these models.

<sup>20</sup> Reason, 1997, p. 10. However, that doesn't remove individual responsibility and accountability for egregious or negligent actions. Reason introduces the concept of a just culture, as distinguished from a blame culture or a no-blame culture, to fairly determine culpability.

<sup>21</sup> Hollnagel, E., Leonhard, J., Licu, T. and Shorrock, S., 2015. "From Safety I to Safety II: A White Paper," Eurocontrol. A key insight that risk managers can also take from the Safety II perspective is much can be learned about risk from observing and analyzing why processes go right rather than focusing exclusively on why