Not Losing Comes First: A Practical Look at the Challenges of Enterprise Risk Management

By Tom Coyne
May 2012

When I was younger, I spent a lot of time and energy on finding ways for my teams to win, without realizing that this was not the same thing as ensuring we didn't lose. Experience finally taught me that the latter is actually the greatest challenge facing most companies. According to the U.S. Small Business Administration, only thirty three percent of firms survive as independent entities for ten years or more. Not losing buys the time that companies almost always need to deliver the big wins they hope to achieve for their investors and employees. Not losing is the essential purpose of the mix of activities that in recent years has come to be known as "enterprise risk management" or "ERM". I've worked on these issues across a range of industries for more than thirty years – first as a banker, then as a consultant, an investment analyst, and most recently as a corporate CFO, CRO, and CEO. Based on that experience, I'd like to offer some practical insights into the real drivers of success for this increasingly critical competence.

The Goals and Value of ERM

Let's start with the goals of ERM, where there seem to be two schools of thought. The first believes that ERM is about optimizing trade-offs between risk and return. Undoubtedly, this is the ideal outcome. However, as a practical matter this approach runs into number of obstacles. The most important of these was first raised by Frank Knight in his 1921 book, Risk, Uncertainty, and Profit. He distinguished between "risky" situations, in which the range of possible outcomes and their associated probabilities were known, and "uncertain" situations, where this was not the case. In his 1954 book, The Foundations of Statistics, Leonard Savage proposed that this difference could be reconciled through the use of subjective probability estimates to convert uncertain situations into ones that are merely risky. However, beyond some

obvious issues (e.g., whose subjective probability estimate to use), this method still falls short if the full range of possible future outcomes is not, or cannot be known in advance (e.g., as is the case with many operational risks, or systemic political or liquidity crises). Moreover, researchers have subsequently found that human beings have a much higher aversion to uncertainty (even when subjective probabilities are used) than they do to risk. In sum, the optimization approach to ERM is hard to convincingly implement when uncertainty predominates, which is usually the case.

However, even situations where we are apparently dealing with risk (e.g., outcomes for which long history or liquid markets exist, like fire losses, commodity prices, interest or exchange rates) present three important challenges to optimization. The first is the potential for problems caused by estimation errors. For example, even when we are dealing with a normal (also known as Gaussian or bell-curve) distribution, the mean is estimated with a standard error. In so far as we as ERM managers are concerned with exposures in the tail(s) of the distribution (i.e., in the range defined by the estimated mean plus or minus two or more standard deviations), these standard errors can have a very large, if often unrecognized, impact on our decisions, if they are solely driven by our quantitative model. The second challenge is the proper form of the distribution to use to describe the range of possible outcomes for a given variable. Mathematically, the Gaussian or normal (i.e., "bell curve") distribution is the most tractable. Unfortunately, research has shown that outcomes produced by complex adaptive systems of interacting human beings (such as commodity or financial markets), as well as non-linear physical processes (like fires and weather) are not Gaussian. Rather, they tend to follow a power law, and have "fatter" and "longer" tails than the normal bell curve. While optimization calculations are easily performed for variables with normal distributions, this isn't the case for variables with power law distributions. To be sure, optimization calculations are still possible, using more sophisticated techniques. Yet even these methods leave a decision maker with irreducible uncertainty about whether the distributions used to describe key variables contain the full range of possible outcomes for them. Unfortunately, backtesting provides no comfort. As Hemez and Ben-Haim have

shown, in highly uncertain situations the better a model is at explaining historical data, the less accurate its predictions of future outcomes are likely to be (Yakov Ben-Haim and Francois Hemez, *Robustness, Fidelity and Prediction-Looseness of Models*, Proceedings of the Royal Society, A, 8 January 2012).

The third challenge you confront when trying to optimize a mix of risky variables is how to describe the relationships between them. In many real world situations, these are very complex, and involve non-linearities and time-delays. Traditionally, the correlation statistic has been used to describe the relationships between key variables; however, it suffers from critical limitations, as correlation only measures the linear relationship between the average levels of two variables. This can lead to substantial underestimates of the true level of risk, particularly in the case of extreme events that are usually of the most interest to risk managers. The limitations of correlation have therefore led to the use of more sophisticated quantitative techniques such as copulas. However, deciding which copula to use to describe the relationship between variables is often made very difficult by sparse or absent historical data about the occurrence of extreme events. Once again, the results are noisy estimates and optimization solutions which both contain an irreducible level of residual uncertainty.

Given the practical problems with implementing the optimization approach, I am in the camp that believes the primary focus of ERM should be on robustness and survival – that is, the search for plans, policies, and processes that will ensure "not losing" under a wide range of possible future scenarios, without sacrificing the possibility of delivering outstanding upside returns. Put differently, in my career I've seen time and again the importance of avoiding failure, which gives a company the time it needs to adapt to changing or unforeseen circumstances in order to deliver the high performance it has promised to various stakeholders. While most business writing focuses on value creation success stories, the data on business survival rates highlights the fact that not losing is the greater challenge for most organizations, though most people don't realize it. With that in mind, I admit to being somewhat ambivalent about the change in terminology used in the new ISO 31000 Risk

Management Standard, which switched the definition of risk from "the chance or probability of loss" to "the effect of uncertainty on objectives." I wholeheartedly support the recognition of Knightian uncertainty as a critical issue; however, I'm not as enthused about weakening the focus on loss and the connection to avoiding failure.

Another issue that I have encountered time and again is that despite the publication in recent years of excellent enterprise risk management frameworks by both COSO and the International Standards Organization, too many business leaders still do not understand its economic value, which usually substantially exceeds its cost.

For example, every investment manager knows that when seeking to achieve a long-term compound return goal it is mathematically more important to avoid large drawdowns than to achieve a few more basis points of return. Or consider a corporate example. Let's say a company that believes the net present value of its assets, projects, and options is worth $1 billion, assuming it survives as an independent entity for ten years. To the extent that survival is explicitly addressed at all in most NPV analyses, it is via the assumption that if the company does not remain independent, it will be because it is purchased at a premium price, rather than that it will fail with substantial destruction of investor value. Let's therefore be charitable, and assume that of the 67% of new companies in the United States that don't survive for at least ten years, 17% are purchased at a premium, and that this outcome is included in the NPV analysis. Let's further assume the remaining 50% of companies that don't last ten years destroy significant shareholder value when they are taken over or close. Applying this factor to the $1 billion NPV leaves a "survival adjusted" NPV of $500 million. Now let's conservatively assume that an effective ERM program can reduce the probability of significant value destruction to only 40%. That adds $100 million to the adjusted NPV. This simplified example makes a key point – as professional investment managers know, the true return on an effective risk management program is usually far greater than most people realize.

I also like to say that while the term is relatively new, there is nothing new about the activities that underlie enterprise risk management. Whether implicitly or explicitly, ERM has always been one of an organization's four core management processes, along with leadership, strategy, and execution. Briefly, leadership defines the purpose and goals of the organization, obtains resources and builds stakeholder relationships, and provides motivation and support to a team. Strategy is about sensemaking and design. The former seeks to identify the key elements in the situation facing a company, how they are related, and how they are likely to evolve in the future. The result is assumptions that form the basis for the design process – determining how to achieve desired goals with available means. Execution implements this strategic design, by developing objectives, metrics, plans, budgets, processes, systems, organization and controls. Finally, risk management ensures survival by identifying and assessing risks and uncertainties, providing warning of adverse changes, mitigating and transferring potential losses, and strengthening organizational resilience and adaptability.

Let's look at some practical aspects of these six facets risk management.

Risk Assessment and Warning

When it comes to risk identification and assessment, the first issue I've encountered is confusion over just what constitutes a "risk." I do not agree with the view that "risks" should be narrowly defined as possible events that would have negative consequences. Rather, I prefer to start with an organization's goals, and then broadly define "risk" as any uncertainty that could affect their achievement. Some of those uncertainties can indeed be described in terms of discrete events; however, others are better described as possible ways in which a situation could gradually evolve, and the consequences of such developments. Think of this approach as a combination of the issues that are typically listed in a Risk Register, along with the "Weaknesses and Threats" that are identified in typical strategy "SWOT" and scenario analyses.

There also seems to be a natural tendency in traditional risk assessments to focus on those risks that are easiest to transfer – i.e., market risks, like interest rate, foreign exchange and commodity price risks, for which there are liquid derivative markets, and traditional loss and liability risks, which can be insured. However, because they can be transferred to others, these risks usually aren't the most dangerous threats to a company's survival. It is therefore critical that the risk assessment process focus sufficient time and effort on developing a deep understanding of strategic assumptions and execution tasks that are critical (in terms of their potential cash flow impact) and highly uncertain (in terms of their range of possible outcomes). Unfortunately, this doesn't always happen, usually because of the political sensitivity of these analyses.

Another issue I've seen in ERM processes is their blind spot with respect to so-called complex risks and failure modes. For example, I've seen companies in process industries that do an excellent job of identifying and analyzing individual risks of human and mechanical failure, but which neglect emergent "socio-technical" risks, such as a situation in which the effects of a mechanical or human failure, interacting with an automated control system, rapidly compound and overwhelm the diagnostic and reactive capabilities of human operators. While system simulators can sometimes help to identify and mitigate such risks, they are still limited by the range of possible scenarios their designers built into them. In point of fact, in complex adaptive systems many risks are emergent, and can't be foreseen. Hence, in my experience, simulation training needs to be complimented with a very aggressive focus on perceiving and routinely seeking to understand the root causes of anomalous system behavior and "near miss" events, to identify and assess the previously hidden complex risks and causal pathways that such events almost always reveal. While anomaly and near-miss analysis is most often associated with operational safety programs, I have found that this is a critical process that has much broader applications in ERM.

Another risk management shortcoming I've frequently encountered is around the issue of model risk. As I define it, a model is the representation of a dynamic situation that helps us focus our attention on key variables, understand causal processes, and

predict future conditions and the potential results of our actions. This definition includes not only models of financial and physical flows, but also project and program plans, which model sequences of activities and decisions, and predict their expected result. As noted above, errors can occur in both the estimation of model input parameters and the form of the model itself. And models that once made accurate predictions can become outmoded if important relationships in the underlying system significantly change over time. For all these reasons, independent verification and validation of models is always a good idea, as are root cause investigations of any significant prediction errors they produce. However, a word of caution is also in order here. Philosophically, it is logically impossible to completely verify and validate a model, as this would require a degree of omniscience that neither the modeler nor the independent assessor can possess. At best, one can say that a model has passed a set of verification and validation tests, and not been falsified by them, which should raise our confidence in its use, while still recognizing that the model itself is a simplified representation of a more complex and uncertain reality. In short, while we can take steps to manage model risk, we can never completely eliminate it.

A final risk assessment issue that is often overlooked is the structure of incentives in the organization. While "high powered" incentives have their virtues, if not carefully designed, they can also inadvertently increase the inherent riskiness of an organization via their impact on behaviors and decision processes.

I've already discussed the challenges of using quantitative tools to assess different risks and uncertainties. This doesn't mean that they shouldn't be employed for risk analysis and aggregation – for example, I am a very strong believer in the use of Monte Carlo simulation for estimating cash flow at risk, identifying combinations of uncertain variables that lead to failure scenarios, allocating risk capital, and framing discussions of risk tolerance. In addition, ISO 31010 provides an excellent overview of various risk assessment tools that an organization can use. However, ERM leaders always need to keep quantitative methods' limitations in mind, and complement them with qualitative tools, like scenario analyses, risk registers and subjective risk

probability/impact matrices. Research has repeatedly shown that a combination of intuitive/experience and data/analysis based approaches to risk assessment is most likely to produce the best result.

Another important area in which I've seen corporate ERM programs fall short is the development of warning indicators and processes. This is in contrast to the public sector, where warning is a critical function within the intelligence community. The CIA's Jack Davis draws a key distinction between tactical and strategic warning: "Tactical warning seeks to detect and deter specific threats -- when, where, and how a specific incident will take place. The objective is to avoid surprise and thus block or blunt damage…Strategic warning addresses perceived dangers in broader terms, in order to inform decision makers on general [risk] preparedness – again to prevent or limit damage…The challenge of strategic warning is to help policymakers decide – in advance of specific indicators of danger – which of the many plausible general threats deserve concerted defensive and preemptive preparations…[Therefore] strategic warning should focus on dangers that are plausible and potentially most damaging, but about the details of which decision makers retain important doubts, and/or which are not considered highly likely to occur" (see *Strategic Warning: If Surprise is Inevitable, What Role for Analysis?* by Jack Davis, published by The Sherman Kent Center for Intelligence Analysis). Enterprise risk management needs to focus on both tactical and strategic warning.

The development of indicators to monitor the evolution of critical risks and uncertainties is a critical part of the warning process. In this area, experience has taught me some important lessons. The first is that certain cognitive biases – in particular, over-optimism (e.g., overestimating the most likely outcome), overconfidence (e.g., underestimating the range of possible outcomes), and confirmation (selectively attending to, and overweighting, new data that supports your existing views) – are extremely powerful and resistant to change, even when you are aware of their existence. For this reason, it is critical to put in place structural processes to offset the potential impact of individual and team biases. For example, I

have found that asking people to specify the one or two pieces of information, which, if they were to appear, would cause them to rethink their assumptions and plans, is a very useful approach, which directly generates warning indicators to monitor. Equally useful is asking someone how they think a third party would assess a situation, and what additional information they would seek. This shift in perspective often reduces cognitive bias. From a Bayesian perspective, critical pieces of information should have two essential qualities – they should be very unlikely to be observed if the current assumptions are true, or if the current plan is working (i.e., they should have high diagnostic value); and they should be obtainable from sources likely to be judged highly credible and reliable by decision makers.

The second lesson I've learned is that complex adaptive systems, such as companies, industries, and economies, are full of multiplicative (rather than linear) processes in which important changes take longer to appear than you expect, but then happen faster than you predict, producing the familiar "S-Curve" phenomenon. Perhaps because the human mind naturally thinks in terms of linear patterns, these multiplicative processes have been a rich source of surprises in my career.

The third lesson is that financial metrics are often inferior warning indicators, as their reporting tends to lag the occurrence of critical changes in the competitive environment ad company operating results. As my longtime controller liked to say, "if we wait for a problem to show up in the financials, it may be too late to solve it."

Finally, in addition to monitoring warning indicators associated with specific risks and uncertainties, it is also important to monitor system level indicators, despite the fact that these are inevitably more ambiguous. The starting point for thinking about these indicators is the dynamics of a complex adaptive system, such as a company. Broadly, such a system exists in one of three states: a chaotic state, an excessively stable state, and a state in the region between the two, where the system is maximally robust and adaptive. Obviously, threats to survival are higher in both the chaotic and excessively stable states, where the system is either over or under-reacting to

changes in its external environment. This theoretical view gives rise to some very practical system level early warning indicators, including rising numbers of "fire drills" generally, and, more specifically, rising levels of operational "near misses" and errors of commission (which in my experience are indicators of approaching the chaotic state), and rising levels of strategic surprise, and errors of omission (which are indicators associated with the excessively stable state). On balance, I have found that the relative importance of these indicators changes over time, with "near misses" more important when a company is young, and strategic surprise more important as it matures.

Mitigation and Transfer

It goes without saying that establishing clear responsibility for mitigating identified risks, along with "getting the basics right" with respect to training, compliance, controls and reporting, are all important aspects of effectively managing an organization's risk exposures by either reducing their probability of occurrence, or reducing their negative impact if they do occur. Yet the implementation of all of these mitigation measures can suffer if team members believe they lack significant benefits, and are imposed on them by external authority. For this reason, I have learned that, to the extent possible, successful mitigation should focus on building risk management tools into teams' existing business processes, in ways that they perceive add value to their efforts. One of the reasons I so strongly support the new ISO 31000 ERM Standard is because of its very strong focus on just this point. In my experience, two of the most effective of these risk tools are "pre-mortems" and "after action reviews." The former adds a critical step to every planning process. When a draft plan has initially has been agreed, ask every team member to write down the answer to this question: "Assume our plan has failed. Why did this happen? What signs did we miss? What should we have done differently?" The resulting discussion almost always substantially improves the final plan. The "After Action Review" is a technique that has been used by the military for years, to learn from their experiences. The key to its success has been a focus on team learning, rather than "blaming and shaming", and reporting. While this

represents a challenge for some corporate cultures, the military has demonstrated the substantial benefits of institutionalizing this process.

These are examples of how I have seen effective mitigation work in practice. Each risk has an identified owner responsible for mitigation, who is supported by a web of first class risk management processes and systems. In addition, enterprise risk management staff not only serve as centers of excellence to support the mitigation efforts of risk owners, but also act as aggregators of risk mitigation plans, to ensure that they are efficient and, critically, that residual exposures and uncertainties are clearly communicated to the players involved in the risk transfer and resiliency planning processes.

Quite frankly, risk transfer driven by talented insurance and banking professionals already works reasonably well in most companies. To be sure, there are always opportunities to evaluate the design and cost effectiveness of specific coverages, to review claims management processes, and to examine the ways in which interest rate, foreign exchange, commodity price and other derivatives are being used and the associated risks modeled and managed. On balance, however, my sense has always been that relatively too much time is spent on these issues relative to less well understood and likely more fundamentally dangerous uncertainties. That said, one exception to this rule that I have seen come up over and over again is the confusing web of relationships between the language used in vendor and other contracts, and the language used in insurance coverages. In my experience, this is a fast changing area where exposures can easily fall through the cracks, and where the support of risk brokers and legal counsel is often critical.

I have also learned to take advantage of the free “loss prevention” services offered by many insurance carriers and brokers. Time and again, they have proven that when it comes to risk mitigation, nobody has a monopoly on insight. Similarly, I have realized over the years that there is more to risk transfer than one initially thinks. While less experienced risk managers tend to focus on hedging and insurance, more

experienced risk managers also focus on how well written contracts and time itself (e.g., delaying decisions) can be used to transfer to other parties the adverse effects of uncertain events. In this case, the key risk transfer challenge is building effective working relationships between risk management and teams who have ownership of the contracting process.

Resilience and Adaptability

As I noted above, a defining feature of a complex adaptive system – such as the competitive and economic environment in which companies operate – is that it is extremely difficult to accurately predict its future behavior. Surprises are inevitable; the challenge is to avoid quickly losing as a result of their initial impact. Resilience is the short-term ability of a team to survive the adverse consequences of a surprising negative change in its environment. In contrast, adaptability refers to an organization's ability to evolve over the medium term, so that changes in its environment do not result in potentially threatening reductions in its effectiveness and efficiency. Both resilience and adaptability are critical to survival.

Most often in ERM, we tend to think of resiliency in physical terms – e.g., building redundant facilities and supply chains, backing up data, and regularly delivering safety training. Clearly, these are all important. However, resiliency also has two other dimensions that as ERM leaders we sometimes overlook. The first is financial, where there is very clearly a trade-off between efficiency (which takes a dim view of slack resources), and robustness (which sees slack as an important buffer against uncertain shocks). The second is psychological. For both individuals and teams, research has shown that psychological resilience to sharp spikes in stress is increased by a shared sense of purpose, by a strong network of empathetic social ties, and by leaders who frame sudden changes as challenges rather than threats – in sum, by a healthy corporate culture. Contingency planning is also a key element of resiliency. For example, this is one reason some private sector organizations have formed and trained specialized incident management teams (modeled on the FEMA standard

widely employed by similar public sector teams), to maximize organizational resiliency and responsiveness if disaster strikes.

Adaptability is critical to organization survival over the medium and long-term – i.e., beyond the horizon of short-term resiliency and contingency planning. There is no doubt that adaptability is on the minds of many C-level executives these days. For example, a recent survey by Deloitte found that "60 percent of CFOs feared their strategies may not adapt well to changing business conditions" (*CFO Signals*, 3Q2011, U.S. CFO Program, Deloitte LLP). And as Bain's Chris Zook and James Allen note in their new book, Repeatability, "the extinction of once great innovators is less often caused by technological or market evolution, and more often by self-inflicted wounds and slow cycles of decision and adaptation." From an ERM perspective, the key issue with respect to adaptability is the health of the organizational processes that give rise to the classic drivers of evolutionary change: variation, selection and retention (implementation).

For example, a work team may seek new ways to better execute existing processes; the business unit of which it is a part may seek to enter new markets, better anticipate customer needs or more quickly respond to competitor offerings; at the corporate level, leaders may seek to improve overall organizational learning effectiveness, or the management of the strategic options portfolio, or to better align internal performance metrics with the success criteria used by investors, customers, and potential employees. All of these are examples of adaptive processes that collectively determine whether an organization will survive and thrive over the medium term. ERM managers need to be concerned with the health of each step in this evolutionary cycle. Key questions to ask include, is an adequate number of sufficiently varied ideas being generated? Are the decision processes used to select ideas for further development reasonably transparent and objective or are they opaque and highly politicized? And is the organization making the multiple changes required to retain and maximize returns from the best ideas, or do too many of them tend to be crushed by the dominant culture?

In my experience, there is no shortage of people who will question whether the health of adaptive processes is really an ERM issue. In response, I can only point to the power law decline in corporate survival rates over time as evidence that the health of a company's adaptive processes, especially in a world of ever faster change, is arguably the most fundamental risk management issue a board and management team must address to avoid losing in the medium and long term.

The Critical Role of Risk Governance and Communication

You simply cannot separate effective communications from effective risk management, which is why this issue plays a prominent role in ISO 31000. Of course, to some extent this is simply stating the obvious, as effective communications are critical to many other aspects of overall organization performance. Yet there are some aspects of this issue that are particular to risk management, including educating employees about the nature of risk, risk tolerance, and risk management processes. Yet these efforts will likely be for naught if an organization fails to address the most critical risk communication issues, which, in my experience, is the effective communication of bad news and differing views. If these aspects of organizational communications are subpar, an enterprise risk management program cannot succeed. Period.

And this inevitably brings us to risk governance. In my experience, the starting point for successfully addressing this issue is facing up to some basic truths about human beings and organizations. To varying degrees, we all suffer from certain cognitive biases that are extremely difficult to overcome, even with training. We also have predictable reactions to loss and to increased uncertainty: both make us fearful, and more likely to hew closely to group opinions, to avoid becoming an outcast at a time of heightened vulnerability. Research has shown that these biases and reactions have lasted so tenaciously over the ages because up to a point they have been adaptive – i.e., they have increased our individual odds of survival. However, beyond a point,

they can also become dysfunctional, increasing the likelihood that effective risk managers who question dominant organizational views will be labeled “excessively negative” and “not team players”. Whether this causes good risk managers to leave, or simply results in their marginalization, the typical result will be a significant increase in the company’s risk profile, too often adopted by default rather than as the result of explicit deliberation and decision by a board.

This highlights three governance imperatives for successful enterprise risk management. The first is that staffing, and the composition of the top leadership team in particular, is a critical driver of long-term ERM success, and a common root cause of ERM failures. If you load up your team with players who score high on over-optimism, overconfidence, and the confirmation bias, and low on their ability to productively handle conflict, you are setting an organization up for an eventual risk management failure. The second governance imperative is that even when you staff a team with reasonably balanced people, you still need effective processes that structure and productively manage the conflicts that are an unavoidable part of ERM’s ultimate success. Simply assigning “risk owners” to each risk is not enough; you also have to directly focus on the way an organization manages (or fails to manage) conflict. On the one hand, there are structural approaches that can improve the management of risk-related conflict, including scenario/stress testing, devil’s advocacy, red teaming, and, as a last resort, ombudsmen. On the other hand, I have also found that it is often effective to include an outside consulting firm as part of your annual risk assessment process, for the very well known reason that it is politically easier for them, rather than inside staff, to challenge beliefs (e.g., about the probability and/or consequences of certain risks) that are strongly held within an organization, particularly among the senior management team. Last but not least, the inescapable nature of these conflicts means that the ongoing support of the board, and the audit committee in particular, for a company’s ERM process is absolutely critical to its effectiveness.

Finding Good Risk Advice

Last but not least, I will briefly touch on another practical risk management challenge I have encountered as both a CFO/CRO and as a board member: whom to turn to for advice. I will preface my observations on this issue with the important note that they are anecdotal, and not meant to be representative of any base rate across companies; unfortunately, if thorough research in this area exists, I have yet to see it. As a practical matter, I found three main sources of outside advice on risk management issues. Our audit firms tended to approach ERM from a control and compliance perspective; moreover, there is always some understandable awkwardness when it comes to discussing some risk issues with your auditor. There was also no shortage of consultants who were willing to undertake specific ERM projects for us, such as setting up our risk register, improving our assessment process and/or systems, or strengthening resiliency in a particular area, such as I.T. or succession planning. Finally, risk brokers were usually quite helpful when it came to assessing traditional risks, and helping us decide on specific transfer strategies for them, whether via insurance or buyer/supplier contracting language. Bankers also provided similar transfer services for risks traded in financial markets, like interest rate, foreign exchange and commodity price exposures. Yet there were other ERM issues that often weighed heavily on my mind, but which did not require a transfer product, a consulting project, or an improvement in our controls and compliance procedures. I often wished that the ERM world had the equivalent of an investment banking, management consulting or law firm partner I could call for practical counsel, to help me think through and address issues that fell outside the bounds of the market's three main risk management offerings. If such firms exist, I was never able to find one, and suspect that I was not alone in my frustration at this gap in current service offerings.

Some Final Thoughts

As I said at the outset, the ultimate goal for risk managers is to enable companies to avoid losing for long enough to allow their plans, options and adaptations to deliver big

wins for investors and employees. In a world of increasing complexity and uncertainty, developing a distinctive competence in enterprise risk management is not only more difficult than most people realize, but also much more valuable.

Yet it is also critical to realize that the avoidance of failure will never simply be a linear result of designing the right processes and procedures; a detailed and regularly updated risk register is unlikely to do the trick. Ultimately, in a complex adaptive system avoiding failure and achieving success are both emergent rather than deterministic outcomes that we will never fully understand, much less control. To be sure, there are many actions organizations can take that will raise the probability of avoiding failure. However, we must take these actions with full awareness that we can never fully understand the complex web of non-linear and time-delayed interactions between them, the behavior of outside parties, and the impact of randomness (or luck, if you prefer). For this reason, no algorithm or simple checklist will ever be able meet the challenge of enterprise risk management; wise and successful leadership in this increasingly critical area will always require a mix of intuition and analytical rigor, courage and humility, and a relentless curiosity and desire to learn.

Tom Coyne consults on enterprise and investment risk management. He began his career as a credit officer at Chase Manhattan Bank. He subsequently spent almost twenty years as a management consultant, specializing in turnarounds and growth. He has also served as the CFO, CRO, and CEO of public and private companies. He is a member of the Professional Risk Managers International Association, and a frequent contributor to investment research publications, helping investors to make sense of uncertainty and reduce their exposure to severe losses. He provided subscribers with advanced warning of both the 2001 technology bubble collapse and the 2008 global financial crisis. He is also a member of the top ranked team in the Intelligence Advanced Research Projects Agency's multiyear forecasting tournament. Tom can be reached at tcoyne@sachuestadvisors.com