

Value-Driven ERM: Making ERM an Engine for Simultaneous Value Creation and Value Protection

Authored By:

John Celona, Strategic Decisions Group
Jeff Driver, Stanford Hospital and Clinics Risk Consulting
Ed Hall, Stanford Hospital and Clinics Risk Consulting

ABSTRACT

Enterprise Risk Management (ERM) began as an effort to integrate the historically disparate “silos” of risk management in organizations. More recently, as recognition has grown of the need to cover the “upside” risks in value creation (financial and otherwise), organizations and practitioners have been searching for the means to do this. Existing tools (such as heat maps and risk registers) are not adequate for this task. Instead, a conceptually new “value-driven” framework is needed to realize the promise of enterprise-wide coverage of all risks: both for value protection and for value creation. The methodology of decision analysis provides the means of capturing systemic, correlated, and value-creation risks on the same basis as value protection risks, and has been integrated into the value-driven approach to ERM described in this paper. Stanford Hospital and Clinics Risk Consulting and Strategic Decisions Group have been working to apply this value-driven ERM at Stanford University Medical Center.

INTRODUCTION

Many organizations recognize the need for ERM and many have programs in place. But, as organizations proceed with implementing one of the many varieties of ERM on offer, persistent questions keep arising. How

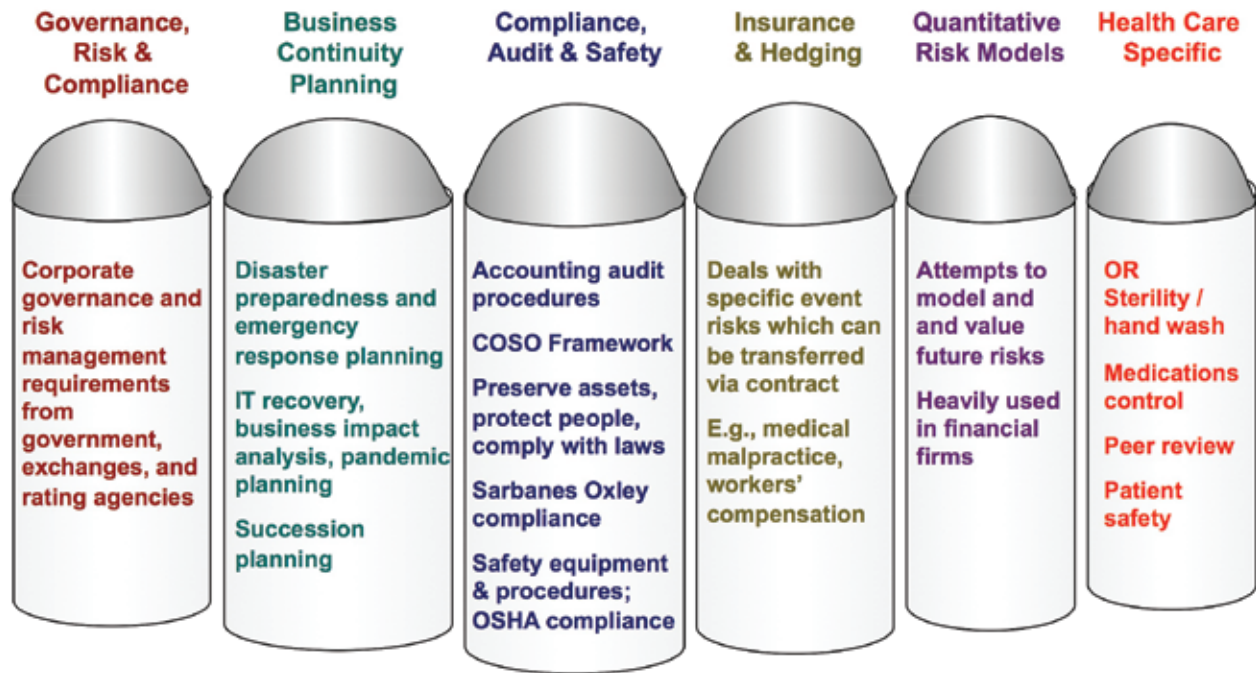
to handle positive opportunities in addition to negative risks? How to link ERM to strategic planning and corporate development efforts? How can we be assured that the risk management being put in place will actually work?

In healthcare, these questions are joined by another daunting array. How do we apply ERM principles to the many risk programs competing for funding and head count, be they Employee Safety Programs, Clinical Risk, Governmental Mandates, Patient Safety and Quality, and so on? How do we integrate ERM into how we actually manage and run the organization? And, perhaps the biggest question of all: how should we respond to the major changes coming from healthcare reform—many of which are not yet clear and are in the process of being defined? These challenges are further compounded by changing demographics (aging population, increasing metabolic syndrome) so getting one’s arms around all the risk and uncertainty seems a Herculean task.

Stanford Hospital and Clinics Risk Consulting, in conjunction with Stanford University, Stanford University Medical Center, and Strategic Decisions Group have been working to apply tested and proven risk management methodology developed over the past 40 years to healthcare. To distinguish these methods from other approaches to ERM, we call it “Value-Driven ERM.”

To understand why a value-driven approach to ERM is becoming a new standard for managing risks and meeting these challenges, we will first briefly recount the history of ERM, then show how a value-driven approach readily meets the challenges that current approaches to ERM may find so daunting.

Figure 1: Historic silos of risk management



ERM: Where It's Been and Where It Is Today

A few decades ago, "Enterprise Risk Management" was not on management's radar. Nonetheless, there were many risk management programs in place to handle specific types of risk:

- Health & safety, which worked in conjunction with OSHA compliance;
- Fire safety, which worked with fire insurers;
- Audit & compliance, which worked with financial controls and accounting auditing procedures;
- Medical malpractice and claims management;
- Regulatory compliance;
- Programs aimed at specific areas of healthcare risk management, such as procedures for operating room sterility and for managing and dispensing medications.

These areas of risk management tended to operate in isolation from one another; they were unconnected "silos" of risk management. Figure 1 shows these historical silos of risk management.

The first step in the development of ERM was simply to re-label all the existing isolated risk management practices as "ERM." Unfortunately, unlike when changing

a drug label, no demonstration of therapeutic effect was required to change the label. As one might expect, the many purveyors of the "re-labeled" varieties of ERM tend to focus heavily on the area they previously provided services in: accounting firms tend to focus on

Figure 2: Basic steps for ERM



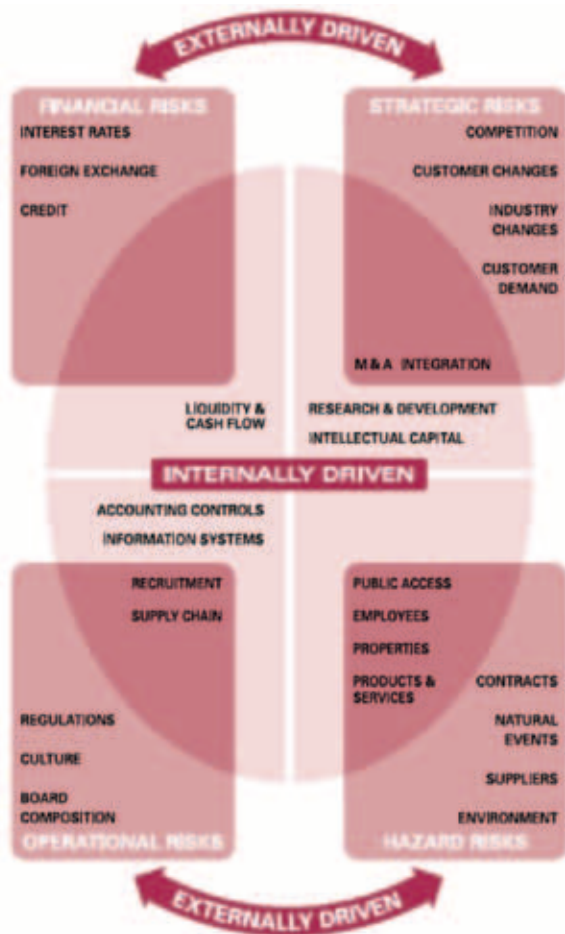
audit and compliance, insurance brokers tend to focus on insurable risk and selling insurance, and so forth.

The second step in ERM's development was to recognize that there were gaps in coverage of risks, and to attempt to fill in those gaps to achieve enterprise-wide coverage of risks. To do this, people early on recognized the need for a process for risk management. The required steps were not too hard to recognize and there is general agreement on what they are:

1. Figure out what the risks are; also called risk identification.
2. Come up with some attributes for the risks that describe their possible effects and allow one to rank them. This is now referred to as risk assessment and evaluation or something similar.
3. Draw up plans for mitigating the most important risks.
4. Monitor your risk mitigation plans for how well they are working.

These basic steps for ERM are illustrated in Figure 2. Of course, key to any process is how one actually executes each step. In particular, how does one identify the risks? And how to assess and evaluate them. The first tools developed were frameworks for identifying risks, heat maps for assessing and evaluating them, and risk registers to show the top risks being mitigated.

Figure 3: Sample framework for eliciting and classifying risks



There are many frameworks for getting people thinking about all the potential risks an organization faces. One sample framework for eliciting and classifying risks is shown in Figure 3.

Initial efforts to rank the relative importance of various risks for the risk register simply by ordering them were not satisfactory. Risk tended to be ranked according to who voiced the strongest opinions or was most senior in the room. In addition, the perceived importance of risk tends to be heavily swayed by what is already in the news – instead of what may be coming in the near or distant future. Two years ago risk registers and mitigation plans were fixated on how to deal with an avian flu pandemic (which never happened!). Today, dealing with an economic meltdown (which already has happened!) is on the top of most lists.

Heat maps were developed to try to lend some discipline to the ranking of risks. The simple heat maps first developed attempted to categorize risks according to their likelihood and impact. A simple heat map is shown in Figure 4.

The two axes will typically have some sort of text descriptions associated with them, with likelihood ranging from “almost never” to “almost certain” and impact ranging from “insignificant” to “critical.” Identified risks are then ranked by placing them somewhere on the heat map.

Once the risks have been ranked, they go on the risk register in order of priority. The prioritized risk register then serves as a launching point for developing mitigation plans for key risks.

Most implementations of ERM today include the “traditional” areas of risk management (the historic “silos”), a process description, a risk register and heat map. Unfortunately, this is not enough to deal with all the risks an enterprise faces.

WHY A VALUE-DRIVEN APPROACH TO ERM IS REQUIRED

Many organizations and ERM practitioners are coming to realize that they are still missing a major portion of the puzzle: the “upside” of risk. How do they deal with the uncertainty around the value creation opportunities an organization has? For healthcare, increasing patient and employee satisfaction, reducing costs and increasing revenues are all upside opportunities. How does ERM deal with these?

In most organizations, the answer is: it doesn't. The corporate strategy and development people have responsibility for growing the value the organization creates for all its stakeholders (staff, patients, physicians, community, and shareholders or trustees). The people responsible for building future success are distinct from those dealing in health and safety, audit, compliance, risk management, insurance, etc.

Let's be clear: all the "traditional" risk management functions are absolutely still required. We sometimes call them "good hygiene." But brushing and flossing and wearing sunscreen will not help you manage your career. Something more is needed.

Pioneering healthcare risk managers are likewise looking for how to apply risk management to increasing value—not just managing downside risks from possible bad events. Value-driven ERM offers a proven way to do this.

There are four major pieces to implementing a value-driven approach to ERM:

1. Take a top-down, value-driven perspective.
2. Focus heat maps on event risks.
3. Implement decision analysis for dealing with

systemic or correlated risks, and for value creation uncertainties and opportunities.

4. Implement a Quantified Risk Appetite (QRA).

TAKING A TOP-DOWN, VALUE-DRIVEN PERSPECTIVE

A value-driven approach to ERM starts with a basic question: what is the enterprise worth? That begs another question: to whom? Effective value-driven ERM must include all perspectives on a firm's risk and value, as illustrated in Figure 5.

At the top, the board of directors is charged with managing the intrinsic value of the organization produced by its operations. For public companies, the expectations and guidance provided by the firm guide its market value. For governmental agencies, we substitute the stakeholder groups for the market value: regulators, customers, employees, community, etc.

The operating business units manage the current value-producing activities of the enterprise, and within them reside the uncertain drivers of value which we are after. Corporate risk management is typically charged with

Figure 4: Simple heat map


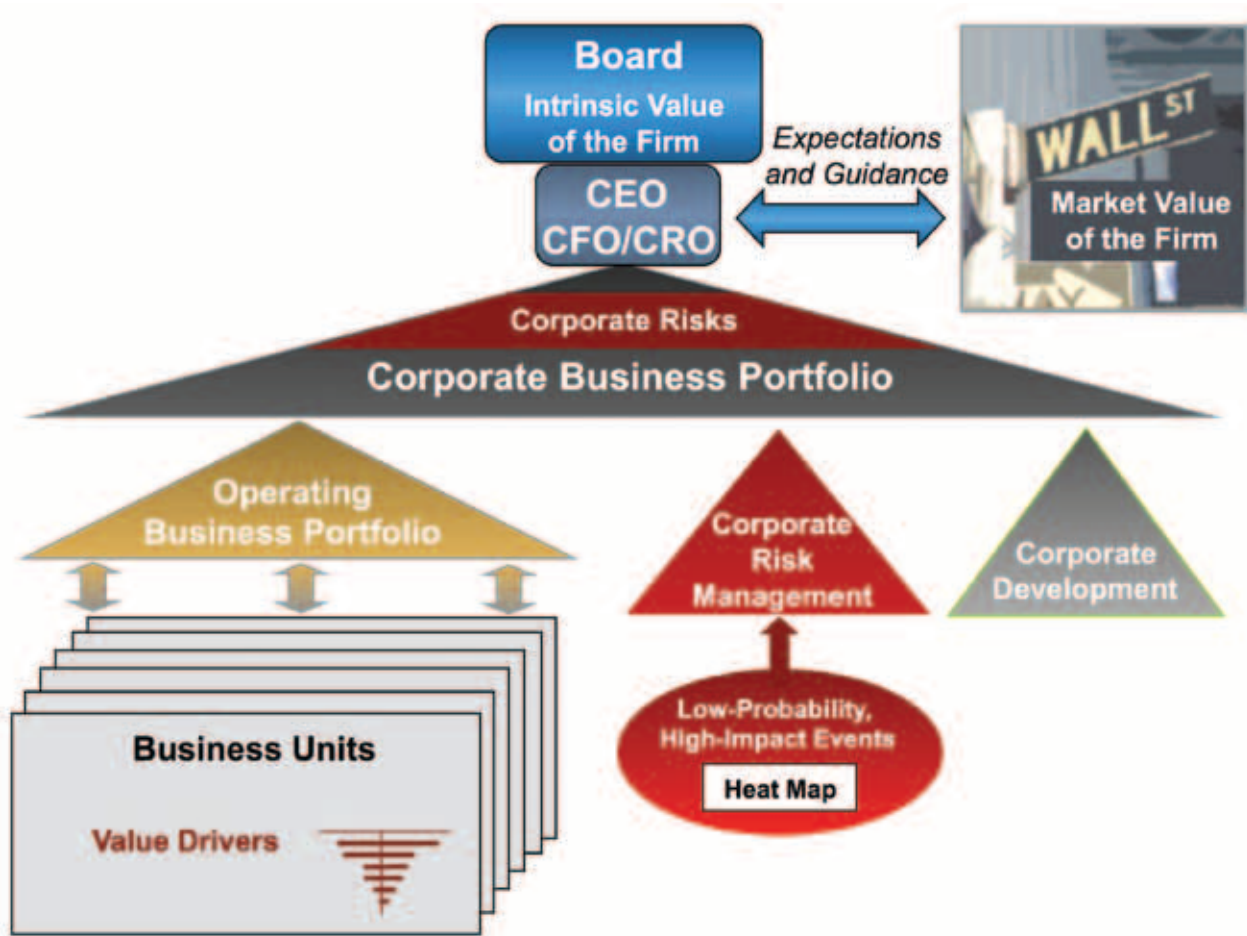


Figure 5: Perspectives and Components of an Enterprise’s Risk and Value



protecting against low-probability, high-impact events (fire, flood, insolvency, regulatory penalties, etc.). Corporate development is responsible for the initiatives that allow the enterprise to continue developing value in the future.

When we look at all the areas of risk and value-creation in an organization, it becomes clear that the answer to the question “what is the enterprise worth?” is itself uncertain. The answer depends on what is happening today and future uncertainties (both risks and value creation). Treating upside uncertainties and downside risks in the same manner allows a unifying perspective: what we are after is the total uncertainty in value. This can be shown in a probability distribution, as illustrated in Figure 6.

The primary job of the board and CEO, then, is to increase the total value (shifting the distribution to the right), and to reduce the risk (narrowing the distribution).

In the extreme low-value situations (the lower left of the distribution) where a distressed organization is sliding towards insolvency and control is going to creditors, the picture of the organization’s value changes dramatically as it deals with financial safety risk. These relationships are shown in Figure 7.

When we look at the total uncertainty in value this way, it becomes apparent that managing the downside risks and upside value creation uncertainties falls into different functional areas of the organization, as illustrated in Figure 8.

Adopting a top-down, value-driven perspective allows redefinition of the problem of achieving enterprise-wide risk management to naturally include upside value creation uncertainty along with downside risk protection. To implement it, it is helpful to separate the process of designing the ERM solutions from operating them on an ongoing basis, as illustrated in Figure 9.

Figure 6: Uncertainty in total value shown as a probability distribution



Once an organization has adopted a top-down, value-driven approach, the major pieces are focusing heat maps on isolated event risks, implementing the tools need for value-creation opportunities and for systemic or correlated risks, and adopting a quantified risk appetite (QRA). We turn to these next.

FOCUS HEAT MAPS ON EVENT RISKS

With the growing recognition that ERM needs to cover the upside of risk to include uncertainty in both value protection (the traditional focus) and value creation, organizations turned to the main tool at their disposal: the heat map. It is common these days to see risks on heat maps like “economic risk” or “competition” or “healthcare reform.”

These risks nestle in the ambiguity present in the simple heat map shown in Figure 3. Just how “likely” is a risk event to occur? And what exactly is the impact?

To strengthen the use of heat maps, organizations need to:

1. Eliminate the ambiguity.
2. Treat the systemic risks (like healthcare reform) using appropriate tools (which are the same as the ones for value creation uncertainties and opportunities).

As a bonus for doing so, organizations will also vastly improve their risk mitigation planning.

Eliminating Ambiguity

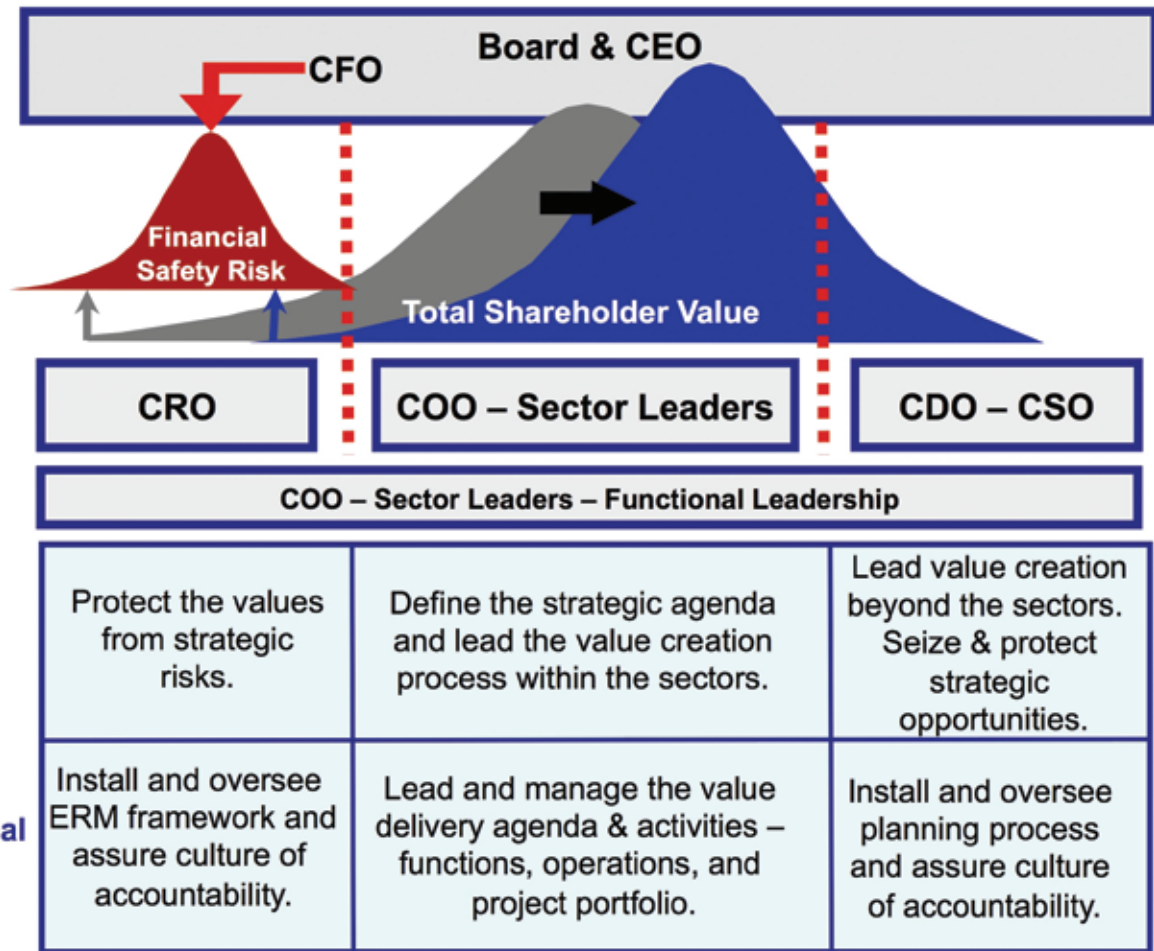
One heat map we reviewed (prepared by a major firm) had the following labels on the likelihood axis:

- Almost certain
- Likely
- Probable
- Unlikely
- Almost never

Figure 7: The primary job of the CEO and the Board is to increase the value, reduce the risk, and manage financial safety risk.



Figure 8: Uncertainty management requires managers in different functional roles to handle specific strategic and operational risks .



What is the difference between “likely” and “probable”? Numerous studies have shown that, when asked to assign a range of probabilities to these terms (e.g., 60-80% likely) people usually assign very different numbers to these terms. This is an exercise we frequently run with groups – with all too predictable results. The same problem arises when characterizing impact with terms like “moderate” and “major.”

The solution is straightforward: to characterize risks with unambiguous numbers instead of vague words. When placing risks on a heat map, we ask for two numbers:

- The probability that a risk event will occur
- The impact of the risk event if it occurs in dollars

This unambiguous form of a heat map is shown in Figure 10.

The \$100 million figure on the impact scale is for illustration and can be adjusted to suit the scale of the risks being evaluated. Using this unambiguous form of the heat map solves the vagueness problem in three ways. First, it removes the ambiguity inherent in using words. People may disagree about what “likely” means, but 60% is 60%—to anyone.

Second, it forces people to think harder about the risk they’re examining. Does a risk event have a 40% or 80% chance of occurring? Over what time frame? What exactly would the impact be if it occurs? How do we quantify that? Disagreements in the answers to these questions usually lead to major insights into how a risk event could play out.

Better Risk Mitigation Planning

The clarity of using numbers also allows better risk mitigation planning. Would a risk mitigation plan reduce the probability of occurrence, the impact if the event occurs, or both? Suppose the risk of a catastrophic fire at a building is thought to be 5% (based on history, the age and condition of the building, the activities in it, etc.). That works out to one of these fires occurring every 20 years. (Does that sound right? Is it too high?) Suppose the mean damage estimate if one occurs is \$20 million (would be higher if we include possible loss of life). This works out to an annual risk exposure of $0.05 \times 20 \text{ million} = \1 million . A fire insurance policy with a \$50,000 annual premium to protect against this \$1 million/year risk exposure sounds like a good deal.

But we're not done yet. Suppose it would cost \$5 million to install an upgraded fire suppression system in the building, and that this would reduce the chance of a catastrophic fire to 1%. The value of implementing this risk mitigation measure to reduce the annual risk by 4% is $.04 \times 20 \text{ million} = \$800,000$. It would take $\$5 \text{ million} \div \$800,000/\text{year} = 6.25 \text{ years}$ for this invest-

ment to pay back in reduced risk exposure. Likely this payback period can be shortened if we include the probable reduction in the annual insurance premium.

On the other hand, a risk mitigation measure (or "risk treatment" as it's sometimes called) is probably not worth undertaking if it's not evident how it would reduce either the probability of occurrence or the impact if the risk event occurs.

This kind of clear thinking about event risks and the potential mitigation or treatment measures for them is impossible until one removes the fog of ambiguity from the heat map and risk event characterization.

Treating Systemic Risks Appropriately

Once the ambiguity in how heat maps are used has been removed, it becomes clear what kinds of risk don't belong on them. What is the probability of "economic risk" occurring? The question makes no sense. Economic risk is out there and manifesting itself in various ways every day. The right questions are what are the prospects for the economy over the time frame you're looking at, and how could that affect your organization?

Figure 9: ERM solution design and ongoing operation

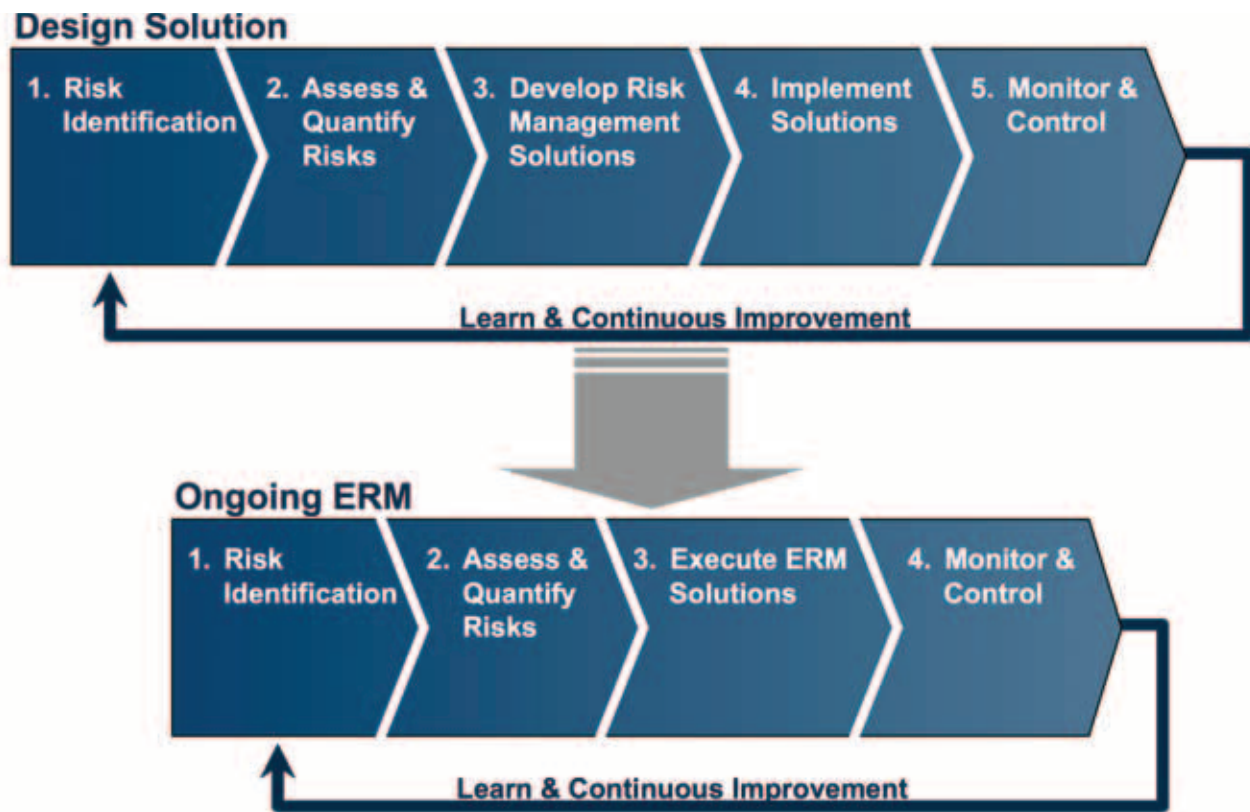



Figure 10: Unambiguous heat map



A simple test reveals which risks should go onto a heat map. Risks should only be displayed on a heat map which are:

1. Isolated event risks (no systemic risks like the economy or healthcare reform); and
2. Risks which are uncorrelated.

The second point cannot be emphasized enough. Risks are often correlated, and the joint effect of one risk leading to another often produces magnified consequences. A large part of the story behind the financial meltdown was missing the correlations. Falling housing prices led to higher mortgage default rates which led to ratings downgrades and collateral calls for the sellers of credit default swaps—and there went AIG.

All other risks (e.g., systemic or correlated risks) need to be handled using methods appropriate for those kinds of risks. They cannot be characterized with just a probability of occurrence and impact if the risk event occurs.

Before delving into those methods, it would be helpful to be clear regarding what we mean by probabilities and risk.

GETTING CLEAR ON PROBABILITIES, STATISTICS AND RISK

“Risk” is defined as “exposure to the chance of injury or loss; a hazard or dangerous chance.” Risks are to be avoided, and are only taken on as unavoidable possible negative side effects of other activities we either really want or are required to undertake. A “risky” deal is a bad one, and taking on risk is often seen as a matter of courage rather than a careful and deliberate consideration of all the prospects.

Trying to redefine “risk” to include the upside attempts to plug the value creation gap in ERM, but seems only to add to the confusion.

Rather than trying to change what people mean by “risk,” we often find it easier to think in terms of uncertainty. The number of patients your facility will treat in 2011 is uncertain. We could define a range of numbers (maybe 20% less than 2010 to 30% more) such that everyone would agree it is highly likely that reality will fall within that range. This serves to define an uncertainty without worrying about whether a 30% increase is a “risk” or not.

In this and our fire example, we have been using “probabilities” to describe the likelihood of events. This

is another source of confusion for people used to dealing with probabilities produced by statistical analysis of data. Perhaps unfortunately, the word “probability” can refer to one or the other of two very different kinds of numbers:

- Statistical probabilities are calculated with statistical analysis of data.
- Subjective (Bayesian) probabilities refer to the opinion of a person (or persons) as to the likelihood of a particular event. These probabilities come from a person rather than a data set.

Certainly, we want people to consider all the relevant data in coming up with their subjective probabilities. We just need to be clear that, when using subjective probabilities, we are using an unambiguous number to represent the opinion of a person rather than words like “highly likely” which different people ascribe different meanings to.

People often initially object to subjective probabilities with questions like “How do you verify the number?” “How do you know it’s the correct number?” A subjective

probability is correct when it correctly reflects the considered opinion of the person it comes from.

Obtaining a correct subjective probability takes some doing. Many papers have been written and a Nobel prize awarded to people mapping out how people intuitively go awry when thinking about uncertainty. Behaviorists have identified over 200 biases in the way people think about and judge uncertainty. Some of the most common ones are sorted into categories and illustrated in Figure 11.

The value in grappling with and using subjective probabilities is that they actually work. In decades of experience using them, we have seen time and again that well-done subjective probabilities do a surprisingly good job in mapping out the possibilities for seemingly intractably complex and uncertain problems. A real example of these methods applied at Stanford University Medical Center follows later in this paper.

Unfortunately, people trained in statistical methods have a very hard time wrapping their heads around

Figure 11: Some common biases when considering uncertainty

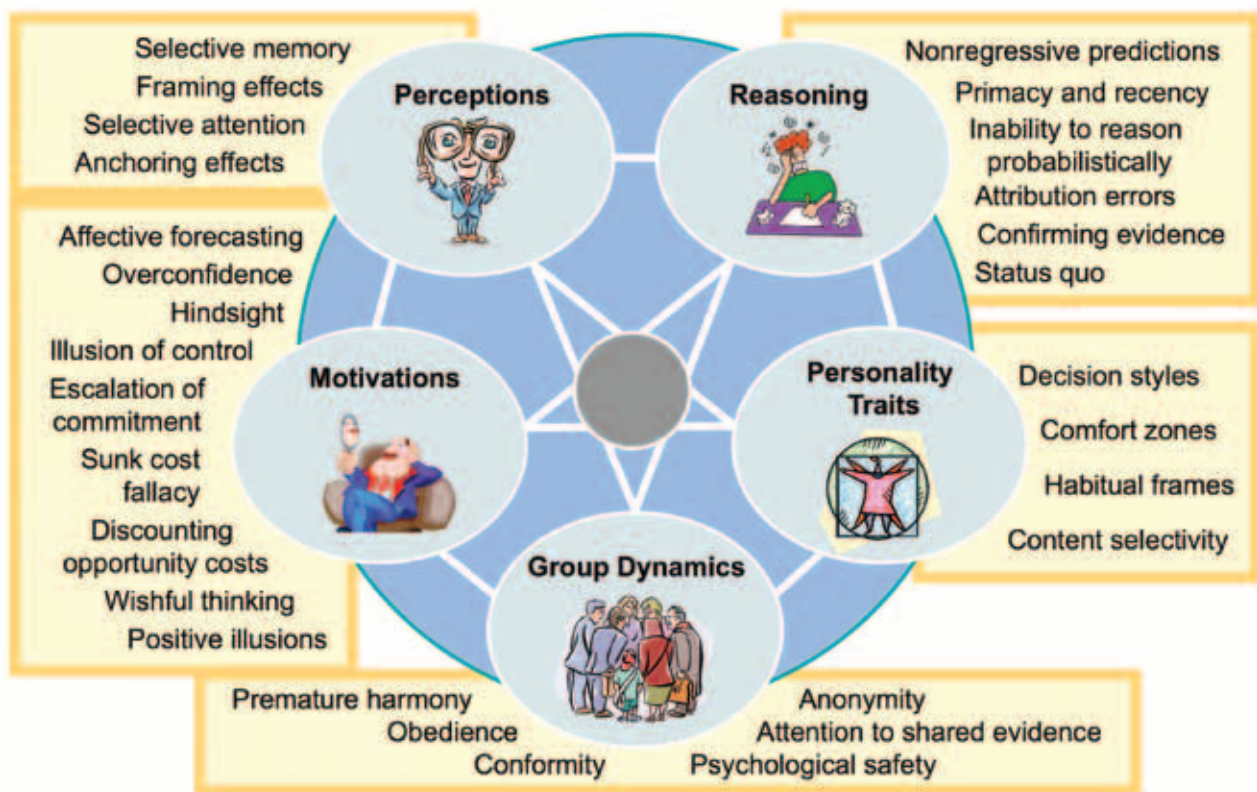


Figure 12: Appropriate use of subjective and statistical probability methods

Methodology	Description	Applications
Subjective Probability Assessment	Careful use of appropriate methods to assess subjective estimates for uncertainties Uses probabilities to describe the likelihood for either discrete events (a fire) or continuous ones (the number of patients next year)	Most risks (competitor actions, capital project cost and schedule, economic variables, etc.)
Statistical Data Analysis	Analyze historical data and construct statistical model for making projections Produces probabilities and significance as outputs of analysis	Only appropriate with a stable system and statistically significant data set

subjective probabilities. They keep looking for a validity test—the way one can test for statistical significance in a clinical or epidemiological study. Unfortunately, many risk management decisions are not susceptible to statistical analysis. You either do something and get an N of 1, or not and N=0. Not much statistical significance there.

For some areas (like life insurance morbidity tables) statistical analysis works very well. The key requirements for using statistics are (a) a stable system; and (b) a representative data set. If either of these conditions isn't met, one needs to use subjective probabilities. These considerations are summarized in Figure 12.

Financial analysts may object that this table omits mention of market models for predicting future events: implied volatility curves and forward curves and the like. Unfortunately, those models only work for the periods when the relevant markets are stable and miss the sudden shifts markets painfully and frequently tend to make. The financial models for predicting default rates on subprime mortgages only worked as long as housing prices were rising (because borrowers could always refinance). When the housing market changed, those models proved spectacularly wrong.

With this background on subjective (Bayesian) probabilities, we are ready to show how they are applied in the tools and methods of decision analysis to handle systemic or correlated risks—both for value protection and value creation.

DECISION ANALYSIS FOR HANDLING SYSTEMIC/CORRELATED RISKS— BOTH FOR VALUE PROTECTION AND VALUE CREATION

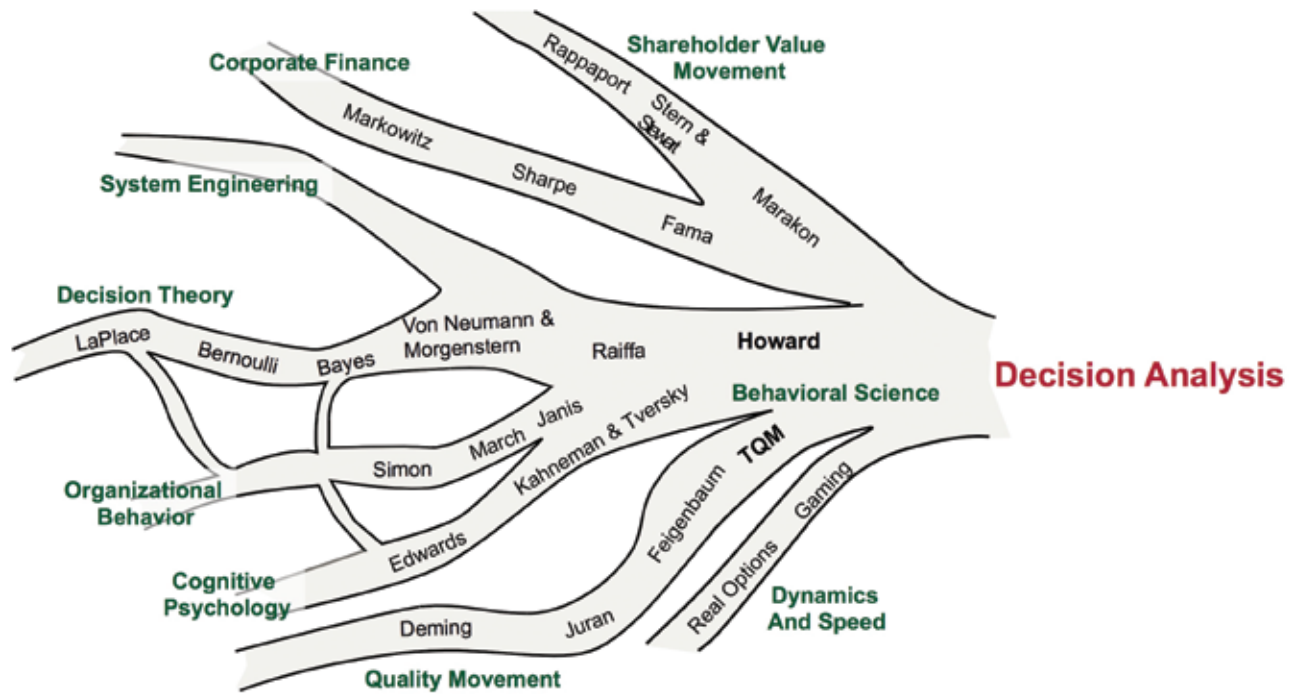
The Origins of Decision Analysis

Decision analysis grew out of efforts to address the challenges of making high-quality decisions under uncertainty. It grew out of the confluence of a number of disciplines, as illustrated in Figure 13.

Learnings from a number of disciplines were drawn together to address the issues arising in making high-quality decisions under uncertainty. Early work in Decision Theory contributed the use of probability to describe uncertainty and how to structure decisions and uncertainties. System Engineering and Dynamics and Speed supplied the means for modeling and analyzing complex decisions and uncertainties and changing dynamics. Cognitive Psychology tackled the problem of how to think correctly about uncertainty, while Organizational Behavior covered decision-making in organizations. Corporate Finance and the Shareholder Value Movement contributed financial metrics and valuation perspectives. The Quality Movement contributed notions of how to evaluate whether a decision is high quality or not. Seminal work in integrating all these threads was done by Dr. Howard Raiffa at Harvard University and Dr. Ronald Howard at Stanford University.

Decision analysis has been extensively applied to medical decisions and in the public policy arena. Central to its approach are identifying, understanding, and quantifying

Figure 13: The Origins of Decision Analysis



all the factors bearing on the costs and benefits of a particular decision. This understanding leads to creation of new alternatives for increasing the value.

Done well, decision analysis produces a robust, transparent, and defensible understanding of the value of various alternatives and a means of identifying how to increase alternative value. This understanding of how specific options create value and the levers for increasing value can be communicated directly to decision-makers without the details of analysis.

For these reasons, decision analysis has become the standard method in a number of industries (including pharmaceuticals) of making investment and program decisions, and of evaluating risks.

The Decision Analysis Approach

Decision analysis applies a “divide and conquer” approach to developing a robust understanding of what the best course of action is and why. A decision is broken down into its component elements, as illustrated in Figure 14:

Alternatives are what one could do. In this context, the decision of interest is deciding whether or not to

implement a safe patient handling program and how to get more value out of the program.

Information and Beliefs include all the information available (historical data) and belief as to what could happen for your organization. For Stanford University Medical Center (SUMC), this included historical data on injuries from moving patients and subjective estimates as to how much those injuries could be reduced with a safe patient handling program. This framework allows for readily updated information and analysis as the program is implemented and results are tracked.

Preferences include a time preference for money (which determines the discount rate for calculating net present value of future cash flows) and a risk preference. Unless the potential consequences are major relative to the total organization (a criteria taken up in more detail later), a quantified risk attitude (QRA) need not be applied.

Logic is captured in the structure of the analysis and, if needed, in a model for calculating values in various scenarios. Because of the complexity, a simple spreadsheet model was required for the safe patient handling program at SUMC.

The Decision is what one decides to do: such as whether to implement a specific risk treatment or value creation program like safe patient handling.

The Outcome is what then happens. It could be whether the risk event occurs, or the results of implementing a safe patient handling program.

A key to assuring robustness in the analysis is using subjective Bayesian probabilities to express judgments on uncertainties, rather than making a single assumption (like on reduction in workers' compensation).

For uncertainties which have a limited number of distinct possible outcomes (like a catastrophic fire occurring or not occurring), we assess the probability of each outcome.

Other uncertainties, like the reduction in staff injuries for implementing a safe patient handling program, have many possible outcomes. In mathematical terms, these are continuous variables as opposed to discrete variables like a fire or no fire. For these continuous variables, we pose three simple questions to bound the range of uncertainty:

- What's a number low enough that there's only a 10% chance the actual outcome would be lower?
- What's a number high enough that there's only a 10% chance the actual outcome would be higher?
- For what number is there a 50/50 chance that the actual outcome would be higher or lower?

This framework allows us to define what a good decision is: one that's logically consistent with the alternatives, information, and preference one had at the time the decision was made. A good outcome is what one hopes will happen.

This framework also assures a transparent, thorough, and defensible understanding of what the key issues are for a particular decision and the range of possible outcomes.

Robustness is also assured by using an iterative approach to analyzing a particular decision. At each stage, questions may be raised prompting further inquiry. This iterative approach is illustrated in Figure 15.

This iterative approach also provides a stopping point. When the analysis is the best it can be given the current information and provides a clear direction for making the decision (be it a new clinic or risk treatment), it is time to end that cycle. The inherent uncertainty in the prospects is often never resolved: you'll just know what did or didn't happen.

Accordingly, the progress of the analysis is not guided by eliminating the uncertainty (which is usually not possible) but, rather, by completing enough analysis to clarify understanding of the uncertainties in the decision, the levers for increasing value, and the best course of action. The analysis is a tool for conducting "thought experiments" to achieve clarity of understanding and clarity of action. We sometimes describe this as striking a reasonable balance between "extinction by instinct" and "paralysis by analysis."

Figure 14: The Elements of a Decision





Figure 15: The Decision Analysis Cycle



VALUE-DRIVEN ERM APPLIED TO SAFE PATIENT HANDLING

Safe Patient Handling (SPH) programs are an example of a value-protection and value-creation activity where the total value potential can be better understood, managed, and increased using a value-driven approach to ERM. Historically, SPH programs have been justified solely from a value-protection perspective: reducing staff injuries and workers' compensation claims. In applying a value-driven approach at SUMC, we discovered this not only missed the largest components of potential value, it also didn't reveal how to get even greater value out of the program.

To achieve these objectives and create the greatest value for patients, staff, facilities, and communities, we first built a comprehensive understanding of the total costs and benefits for an SPH program (downside risk and value-creation uncertainties). Second, we used the understanding of key value drivers to create new alternatives that increased total program value.

Background

Safe Patient Handling Programs install in facilities a variety of means to assist in lifting, turning, and transporting patients. These assists are becoming more important as decreasing patient mobility places increasing demands on care givers. The benefits of these programs include:

- Reduced patient falls and the costs associated with them;
- Reduced patient ulcers and treatment costs;
- Increased patient satisfaction;
- Increased referrals from satisfied patients;
- Reduced staff injuries;

- Reduced costs from workers' compensation and lost or restricted work days;
- Improved worker satisfaction;
- Improved worker retention and reduced turnover costs.



For some of these benefits, past studies provide data on results achieved in facilities that have implemented SPH programs. However, for all of these benefits (and for the costs, as well), the critical questions are determining what the future benefits from implementing a program would be in a particular facility, the right level of investment, and how to increase total value.

Because we are dealing with future impacts (and because there is uncertainty in those impacts), it was necessary to utilize the tools of decision analysis and subjective (Bayesian) probability for understanding uncertainty in future value and for creating alternatives to increase that value.

Understanding the Value

The first step in the process is identifying and structuring all the factors which need to be considered, including the decision(s). These are represented in an influence diagram. Below is the diagram developed by SDG/Stanford Risk Consulting for the SPH program at Stanford Hospital and Clinics (SHC).

Interpreting the influence diagram is straightforward:

- Decisions are indicated by boxes;
- The ultimate net value appears in a hexagon;
- Uncertainties appear in ovals;
- Arrows indicated the relationships between factors.

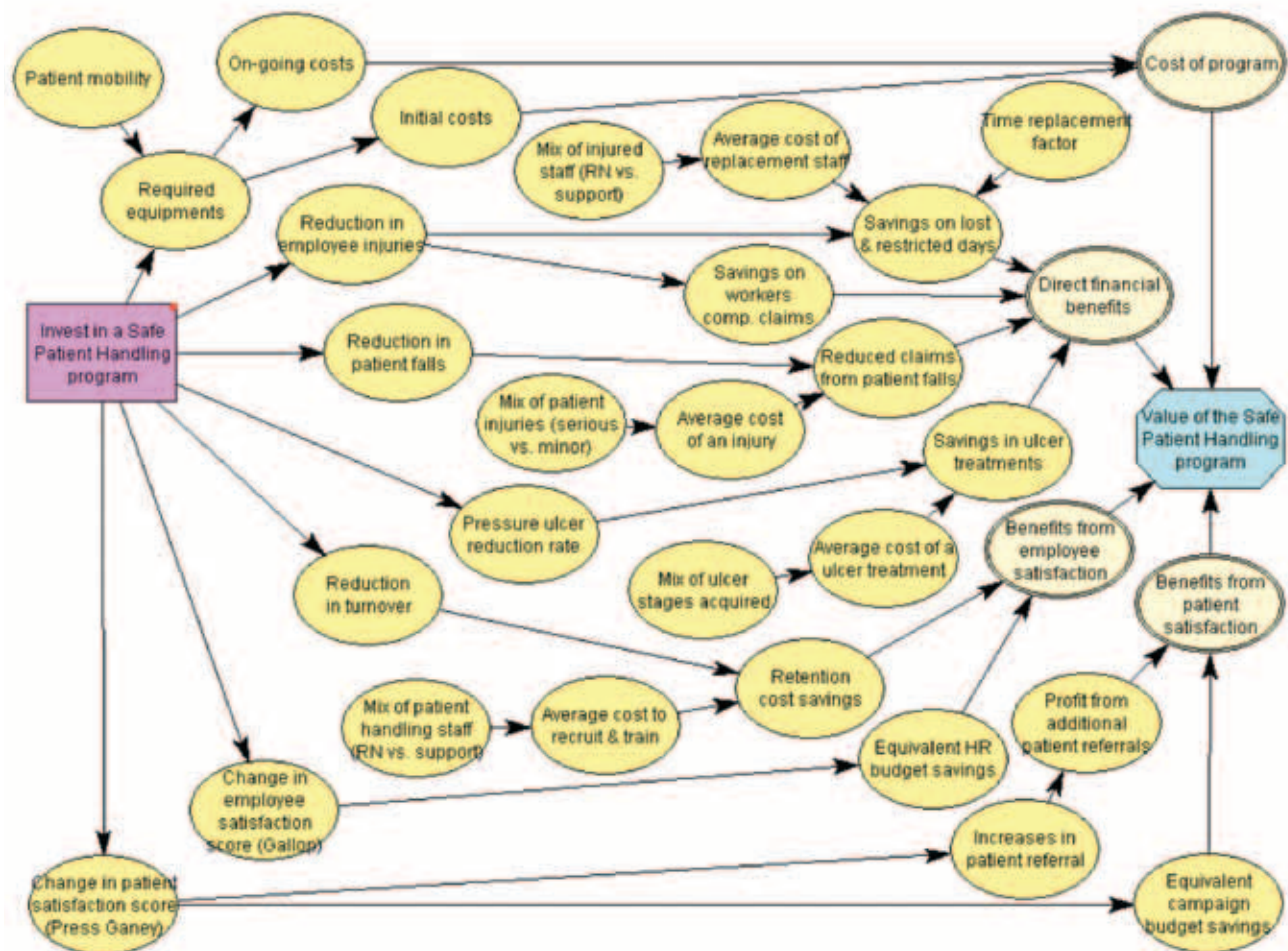
The influence diagram furnishes a map for creating a simple spreadsheet model to calculate the program value. In contrast to usual spreadsheet models which are

only geared to calculating the value in a single scenario and perhaps a few alternates, we need a model which can calculate the total program value in any possible scenario. This is necessary to understand the total uncertainty in program value.

The next step is to create estimates for all the input factors. One example is the reduction in staff turnover from lower physical job demands. In contrast to traditional methods using only one number as an estimate, we use a range of numbers to express the uncertainty in the actual result achieved. Using ranges for inputs achieves three purposes:

1. We can be highly confident that actual results will fall within the range assessed.
2. It enables quantification of factors which are difficult to quantify.

Figure 16: Influence diagram for the total costs and benefits of the SPH program at Stanford University Medical Center



- It enables identification of which factors are the most important drivers for program value.

For example, for the Stanford Hospital and Clinics SPH program we assessed a range of 0–2%–20% for the reduction in turnover. Calculating the program value with each input set to its various values and then arranging these from largest value change to smallest produces the tornado diagram in Figure 17.

Surprisingly, the biggest driver for increasing the value out of SHC’s SPH program was reduced turnover – despite SHC already having a very low turnover of only 2% per year in the affected staff positions. We quantified the value of lowering turnover by calculating the reduction in training costs for new staff. The training cost for a new nurse is a well-studied number and we used an average cost of \$60,000 to recruit and train a new nurse. Similar methods were used to quantify the impacts of all the risk factors.

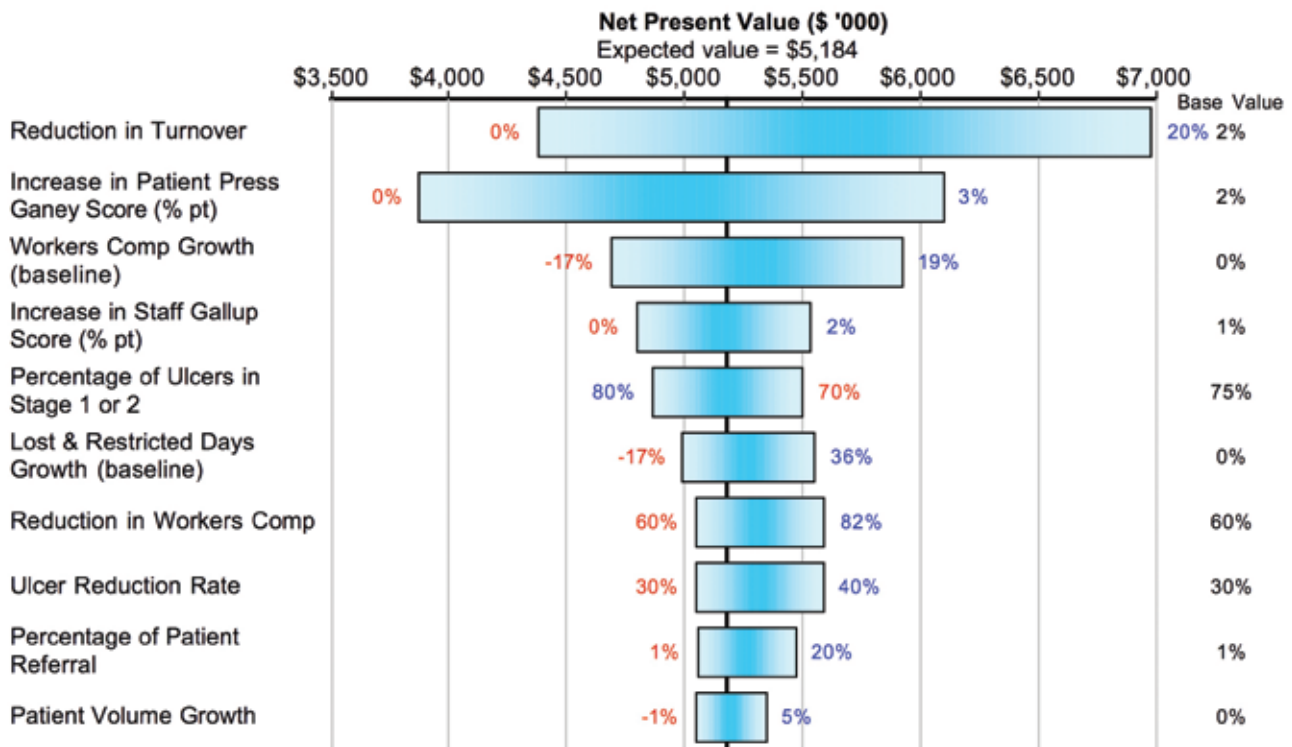
The key value drivers identified by the tornado diagram are candidates for improved estimates, either by consulting other experts or by seeking additional data.

Once we are satisfied that the estimates for critical inputs are the best we can make given the data available now, we can create the probability distribution on the total value of the SPH program at SHC (Figure 18). This probability distribution is one of the pieces which add up to the total uncertainty in the value of the enterprise.

We usually show probability distribution in this form rather than bell curves because they’re easier to read. In the worst case scenario, the SPH program adds \$2 million in value (with all the investment costs figured in). It could add as much as \$12 million in value, and the mean (expected value or EV) contribution is about \$5 million. All these figures are net present values of future costs and benefits.

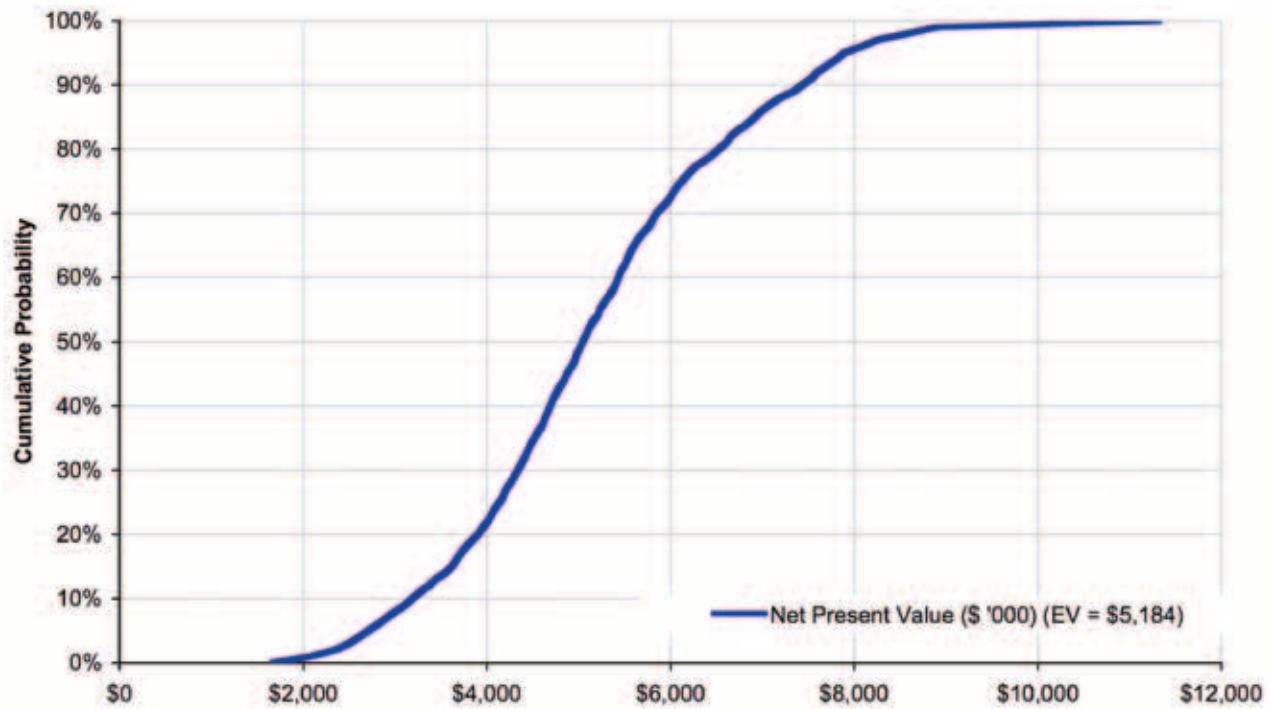
Because there are many components to the benefits of an SPH program, it is helpful to show a breakdown of the costs and benefits by category. We can take the mean contribution of each component and show how they add up to the overall mean value of \$5 million as shown in Figure 19.

Figure 17: Tornado diagram for the value drivers in Stanford Hospital and Clinics SPH program



Note that the tornado diagram is the means of showing the key value drivers illustrated in Figure 16.

Figure 18: Cumulative probability distribution on the value of the SPH program at Stanford Hospital and Clinics.



The ordering of benefits is arbitrary, and they can be easily reordered depending on the priorities of the decision-makers. If a question comes up as to the uncertainty in each value component, we can refer back to the tornado chart.

Increasing the Value

The second and critical part of the process is creating new alternatives to increase program value. The comprehensive understanding of costs and benefits makes this possible.

For example, for the SPH at SHC, change in employee turnover is the single biggest driver of uncertainty in program value. Suppose that investing \$100K in employee communications on the benefits of the program could drive reduction in turnover from a base estimate of 2% to a 10% reduction. From the tornado diagram, we can see that would increase total program value by around \$1 million—a 10-to-1 return.

In contrast, traditional valuation methods only cover cost and benefits that have been verified by extensive studies—in this case, workers' compensation and lost and restricted days. This misses much of the total value

and does not create the understanding for formulating new alternatives to increase the program value.

By applying the tools of decision analysis, we can understand and increase the total value potential in any program—whether safe patient handling or building a new hospital. It easily handles correlated or systemic risks, ensuring, for example, that multiple programs aiming to reduce turnover or increase patient satisfaction produce realistic results when added together. The value-driven approach to ERM integrates at a value level the various programs in a healthcare facility, the same way it integrates the historic silos of risk management.

IMPLEMENTING A QUANTIFIED RISK APPETITE (QRA)

The last major piece in implementing a value-driven approach to ERM is a quantified risk appetite (QRA).

As with risk in general, there is a lot confusion today about “risk appetite” and how to handle it. The basic idea of wanting some way to limit the risks an organization takes is sound. The problem is how to do that consistently and sensibly.

Figure 19: Waterfall chart of contributions to value of the SPH program

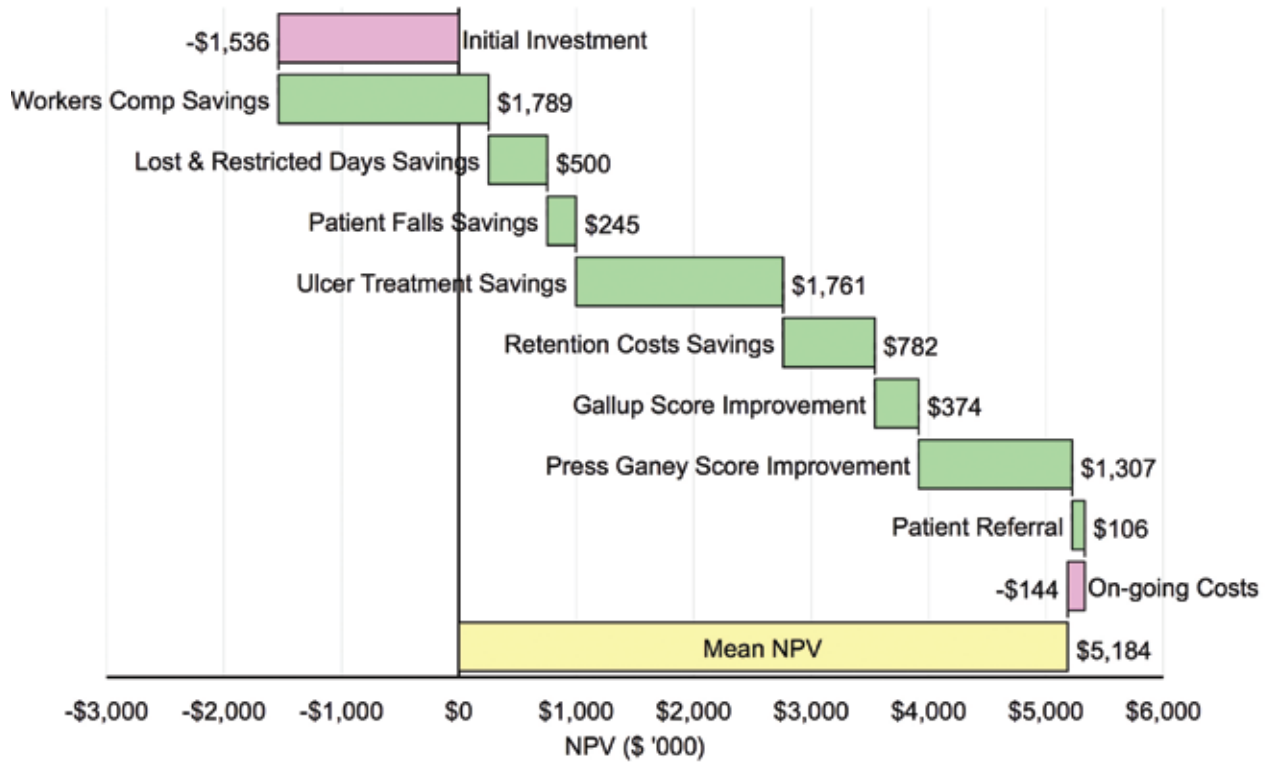
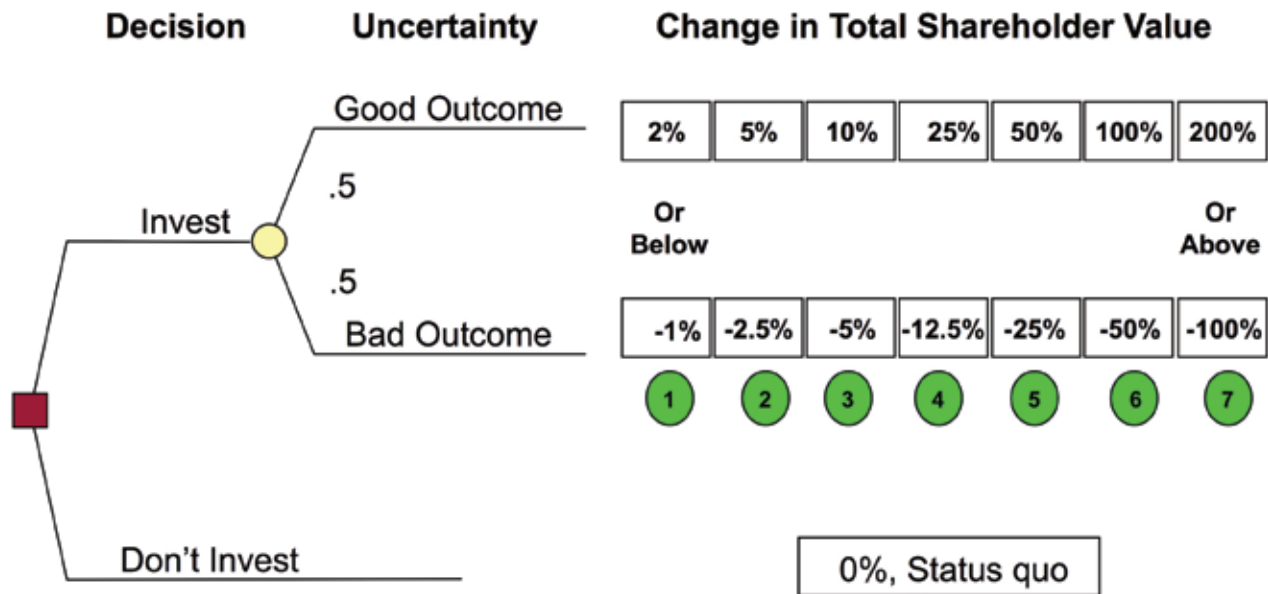


Figure 20: Potential opportunities for assessing a quantified risk appetite



What we commonly see today is fairly arbitrary declarations for an organization that it has, for example, “a low risk appetite for compliance, health, and safety” and a “higher risk appetite for strategic objectives.” Probably what this means (it’s from an actual example) is that bad events in compliance, health and safety have greater negative consequences than missing a strategic target. The question is how to put this into practice.

Other organizations set equally arbitrary limits (e.g., “we won’t tolerate more than X”) without much understanding of why the limit is X rather than X-1 or X+10 or the value consequence of either.

What is needed is an approach to risk appetite that accurately reflects management’s attitude towards taking risk and which can be readily and consistently applied to all risk management activities—whether value protection or value creation or a mix of both.

Helpfully, the most difficult part of achieving this is quantifying the uncertainty in the value of the enterprise (the objective of applying heat maps correctly and value driver/decision analysis tools). Once this is done, applying a quantified risk appetite is very straightforward. It can be applied on large, critical programs (like building a new hospital or buying a hospital system) even if the uncertainty in the rest of the enterprise has not yet been quantified.

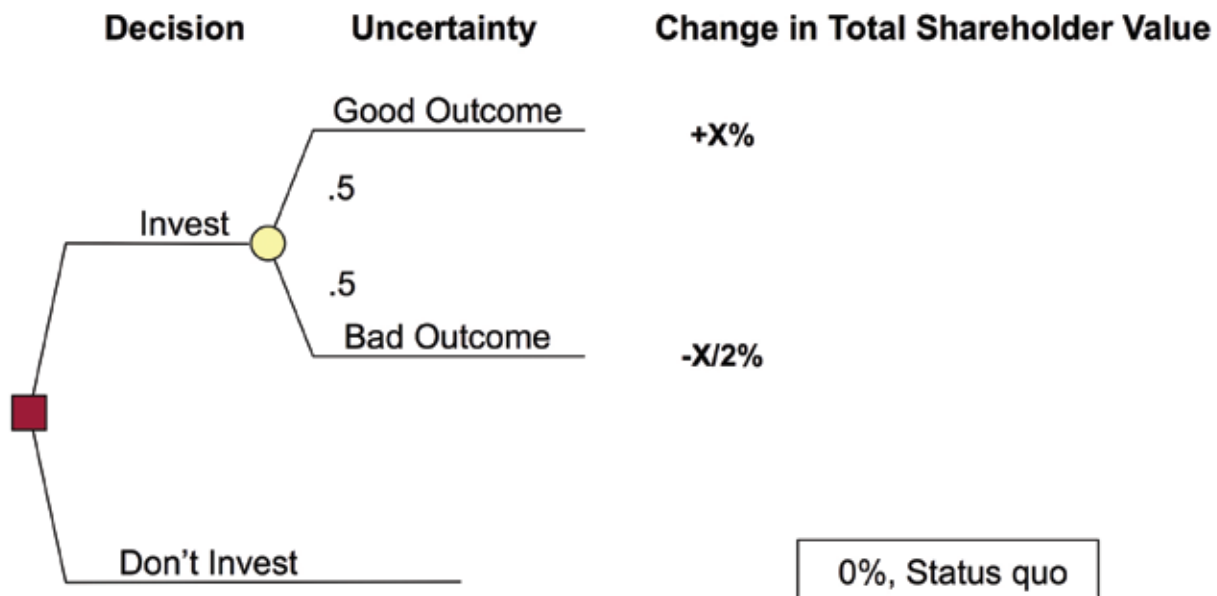
The key is recognizing that risk appetite can be quantified by looking at the maximum risk an organization (or individual) would consider undertaking as characterized by specific opportunities.

Taking the board/CEO perspective, would you invest in an opportunity with a 50% chance of increasing total shareholder value 10%, and a 50% chance of decreasing it by 5% (this a .50 chance of X and .50 chance of -X/2 opportunity). What about increasing shareholder value by 100% or decreasing it by 50%? We can arrange a series of such opportunities and ask what the largest one is the board/CEO would undertake, as illustrated in Figure 20.

The largest possible value for X an organization (or individual) would undertake is the quantified risk appetite (QRA) as illustrated in Figure 21.

The QRA plugs into a simple formula which allows one to calculate the risk-discounted value of any opportunity – the same way a discount rate allows one to calculate the net present value of any series of cash flows. It works like temperature where the scale is arbitrary (water can freeze at 0°C or 32°F). As illustrated in Figure 22, increasing risk aversion (smaller QRA) shows up as a curved line. With risk aversion, negative consequences have much greater negative (risk discounted) value than with a risk neutral, mean value approach. Likewise, big potential payoffs are valued less.

Figure 21: The largest X is the Quantified Risk Appetite (QRA)



The formula (an exponential) converts possible outcomes into utility. Then the quantified uncertainty (already done!) is used to calculate the expected utility, that is then converted back to show what the risk discounted value of any opportunity is.

The great strengths of this approach are that it matches how people actually view risky prospects and that it can be used to consistently calculate a risk discounted value for any prospect – whether risk mitigation or value creation. Ask yourself: would you take a 50% chance of winning \$1 or losing 50 cents? How about winning \$100 or losing \$50? Or winning \$100,000 versus losing \$50,000? It's that simple.

We have been through this exercise to establish the QRA for organizations, but there is an easy shortcut: for most organizations, risking between 10% and 25% of total shareholder value is about all they're willing to take on.

One consequence that becomes clear from this approach is that, for most of the decisions companies make, they don't need to consider risk appetite at all! The possible consequences are too small to require taking risk appetite into consideration. It's only the really big opportunities (or risks) where it comes into play. Organizations have an immediate opportunity to increase value creation by eliminating unnecessarily risk averse decisions at lower levels. Typically, managers have risk aversion relative to their budgets – or their compensation targets. The result is a major value gap, as illustrated in Figure 23.

CONCLUSION

To achieve the entire promise of enterprise-wide risk management, organizations need to adopt a value-driven approach to ERM. This allows all value-protection and value-creation opportunities to be evaluated on the

Figure 22: The QRA shows the degree of risk aversion

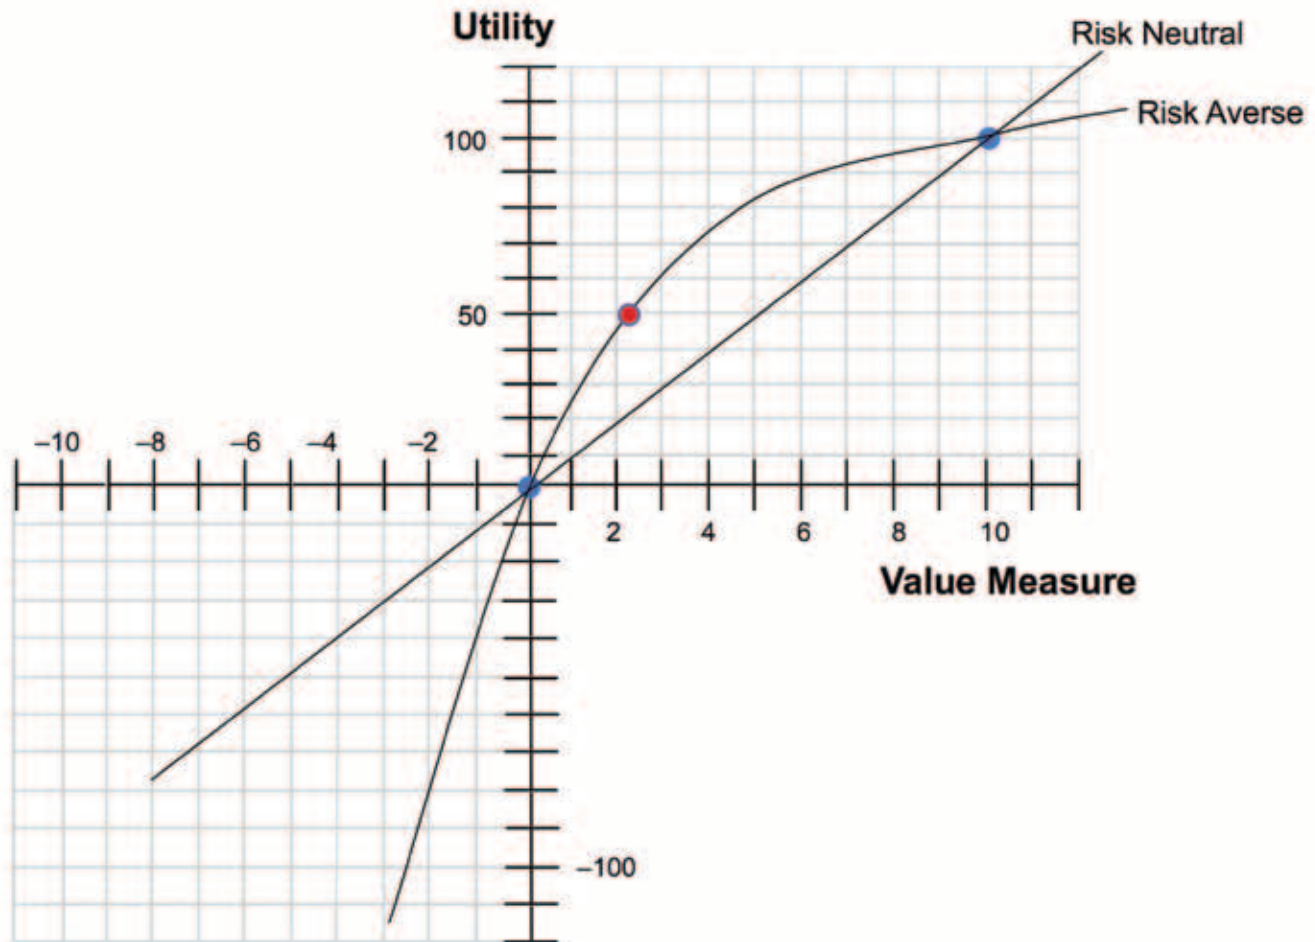
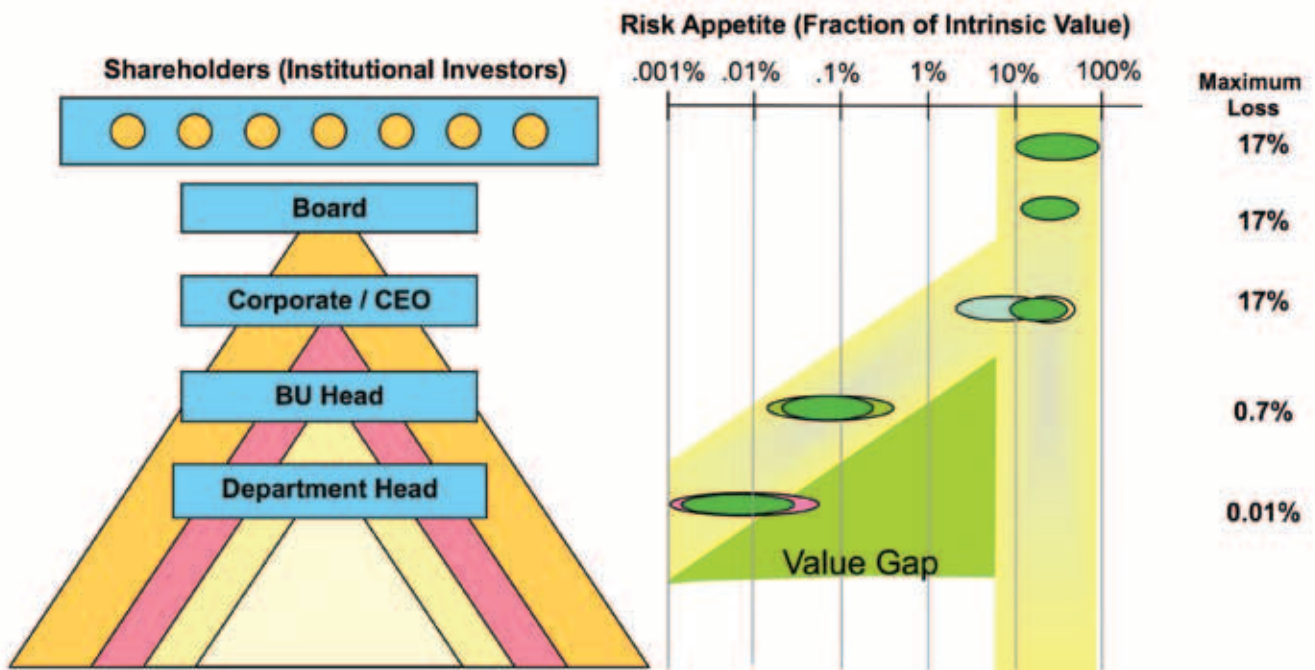


Figure 23: Value gap from inconsistent attitudes towards risk.



same basis, improves risk management, and unlocks the value drivers for increasing total value.

There are four major pieces to implementing a value-driven approach to ERM:

1. Take a top-down, value-driven perspective
2. Focus heat maps on event risks
3. Implement decision analysis for dealing with systemic or correlated risks, and for value creation uncertainties and opportunities
4. Implement a Quantified Risk Appetite (QRA).

Cost of risk can be reduced, and opportunity-value capture increased, using an enterprise value ERM perspective. An enterprise’s largest exposure: business model risks, becomes everybody’s business. This captures the upside of risk by including business model risks in the system.

Cost of risk is reduced by better understanding of the drivers of risk, and of the costs and benefits of mitigation plans. Cross-functional exposures are uncovered by the systematic quantification of uncertainty.

In addition, existing investment in ERM can be leveraged into an accelerated adoption of a value-driven

approach. Decision-makers then get a clearer picture of their business and are thus positioned to manage more profitably, more effectively, and with less risk.

This approach solves the problems we commonly see in ERM implementations:

- Backwards-looking methodology for future risks
 - Actuarial methods
 - Not considering systematic shifts (not “normal times”) – major discontinuities can be normal
- Ad hoc approach to risk attitude
- Bogging down in an extensive risk inventory
 - Lack of a value-driven perspective: everything is important
- Focus on event risk and missing value drivers
- Missing interconnectedness
- Failure to account for psychological distortions in risk judgment
- Focus on only the downside (i.e., risk is a four-letter word)
- Compliance mentality

References

1. Crisis on Wall Street - <http://blogs.wsj.com/wall-streetcrisis>, posted by WSJ Staff on September 16, 2008.
2. http://en.wikipedia.org/wiki/Daniel_Kahneman
3. Buttonwood: "Heart of darkness," *The Economist*, October 27, 2007, p. 86
4. http://en.wikipedia.org/wiki/Decision_analysis

Additional Readings

Allen, Michael S., *Business Portfolio Management: Valuation, Risk Assessment, and EVA Strategies*, New York: John Wiley & Sons, 2000.

Amran, Martha, and Nalin Kulatilaka, *Real Options: Managing Strategic Investment in an Uncertain World*, Boston: Harvard Business School Press, 1999.

Barabba, Vincent P., *Meeting of the Minds: Creating the Market-Based Enterprise*, Boston: Harvard Business School Press, 1995.

Behn, Robert D., and James W. Vaupel, *Quick Analysis for Busy Decision Makers*, New York: Basic Books, Inc., 1982.

Bodily, S., *Modern Decision Making: A Guide to Modeling with Decision Support Systems*, New York: McGraw-Hill Book Co., 1984.

Brown, Rex V., A. S. Kahn, and C. R. Peterson, *Decision Analysis: An Overview*, New York: Holt, Rinehart, and Winston, 1974.

Bunn, D. W., *Applied Decision Analysis*, New York: McGraw-Hill, 1984.

Clemen, Robert T., *Making Hard Decisions: An Introduction to Decision Analysis*, Duxbury Press, 1996.

Cotovello, Vincent T., and Miley W. Merkhofer, *Risk Assessment Methods: Approaches for Assessing Health and Environmental Risks*, Plenum, 1993.

Driver, Jeffrey F., and Farrokh Alemi, *Forecasting Without Historical Data: Bayesian Probability Models Utilizing Expert Opinions*, *Journal of Medical Systems*, Vol. 19, No. 4, p. 359-374.

Golub, Andrew L., *Decision Analysis: An Integrated Approach*, New York: John Wiley & Sons, 1996.

Hammond, John, Ralph L. Keeney, and Howard Raiffa, *Smart Choices: A Practical Guide to Making Better Decisions*, Boston: Harvard Business School Press, 1998.

Holloway, Charles A., *Decision Making Under Uncertainty: Models and Choices*, Englewood Cliffs, NJ: Prentice-Hall Inc., 1979.

Howard, Ronald A., and James E. Matheson, eds., *Readings on the Principles and Applications of Decision Analysis*, 2 volumes, Menlo Park, California: Strategic Decisions Group, 1984.

Keeney, Ralph L., and Howard Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, New York, NY: Cambridge University Press, 1993.

Keeney, Ralph L., *Value Focused Thinking: A Path to Creative Decisionmaking*, Boston: Harvard University Press, 1996.

Luenberger, David G., *Investment Science*, Oxford University Press, 1997.

Matheson, David, and James Matheson, *The Smart Organization: Creating Value Through Strategic R&D*, Boston: Harvard Business School Press, 1998.

McNamee, Peter and John Celona, *Decision Analysis for the Professional*, fourth edition, Menlo Park: Smart-Org, Inc., 2007.

Merkhofer, M. W., *Decision Science and Social Risk Management*, Boston: Reidel, 1987.

Newendorp, Paul D., *Decision Analysis for Petroleum Exploration*, Planning Press, 1998.

Raiffa, Howard, *Decision Analysis*, McGraw Hill College Division, 1997.

Raiffa, Howard, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, New York: Random House, Inc., 1986.

Samson, Danny, *Managerial Decision Analysis*, Homewood, IL: Richard D. Irwin, Inc., 1988.

Schlaifer, Robert, *Analysis of Decisions under Uncertainty*, Melbourne, FL: Robert E. Krieger Publishing Co., Inc., 1978.

Skinner, David C., *Introduction to Decision Analysis*, Probabilistic Press, 1999.

Spurr, William A., and Charles P. Bonini, *Statistical Analysis for Business Decisions*, Homewood, IL: Richard D. Irwin, Inc., 1973.

Thompson, Mark S., *Decision Analysis for Program Evaluation*, Cambridge, MA: Ballinger Publishing Co., 1982.

Tummala, V. M., and Richard C. Henshaw, eds., *Concepts and Applications of Modern Decision Models*, Michigan State University Press, 1976.

von Winterfeldt, Detlov, and Ward Edwards, *Decision* Cambridge University Press, 1988.

This monograph is part of a series of timely summaries on critical risk management issues presented by the American Society for Healthcare Risk Management. Reproduction for distribution without permission is prohibited. Request permission via e-mail at ashrm@aha.org.

Reprints must include the following information: © 2010 Strategic Decisions Group International, LLC

This material is not to be construed as providing legal advice. Compliance with any of the recommendations contained herein in no way guarantees the fulfillment of your obligations as may be required by any local, state or federal laws. Readers are advised to consult a qualified attorney or other professional on the issues discussed herein.

