January 2011

Managing a portfolio of risks

Jeffrey Keisler
University of Massachusetts Boston, jeff.keisler@umb.edu

Igor Linkov

MANAGING A PORTFOLIO OF RISKS

Jeffrey M. Keisler

Igor I. Linkov

With the assistance of Drew Loney

ABSTRACT

Managers and stakeholders concerned about all sorts of entities often face a complex mixture of threats originating from multiple sources. Traditional management approaches that deal with multiple decisions, such as project portfolio management and risk assessment, are focused on single sources of risk or integrate risks using simplified assumptions (e.g., simple addition). Current challenges – including emergency response to natural disasters, homeland security threats, emerging materials, and correlated events – and an uncertain future with respect to societal concerns – for example climate change, economic development, and social/religious instabilities – require integrated application of these methods. This chapter describes a range of risk approaches that enable a flexible approach to risk portfolio management. Decision quality serves as an integrating framework for applying a mix of analytic methods to facilitate successful risk portfolio management. Methods including portfolio decision analysis and risk management as well as other tools from risk analysis, utility theory, and modern portfolio theory from finance help to address complexities of current risk portfolio management challenges.

1. INTRODUCTION

Risk of harm and loss are present in many aspects of modern life, and as society and technology become more complex, there are ever more risks to manage. This chapter is aimed at readers who already know about managing risk, and describing issues that arise in managing *a portfolio* of risks. The terms risk and portfolio have many definitions, so working definitions are useful. The Society for Risk Analysis gives the following definition of risk: "The potential for realization of unwanted, adverse consequences to human life, health, property, or the environment; estimation of risk is usually based on the expected value of the conditional probability of the event occurring times the consequence of the event given that it has occurred." In this chapter, a single risk is taken to mean a possible event associated with a resulting consequence. A *portfolio* refers to a set of objects that are considered together. For purposes of discussion, there are several distinctions that may be useful. A portfolio of risks is a portfolio of possible events that may result in a portfolio of consequences of concern. There may be a portfolio of at-risk assets, affected by events thereby resulting in consequences. Importantly, the risk manager chooses a portfolio of decisions aimed at reducing risks (and which may contribute to and complicate the portfolio of consequences).

This positioning of the chapter has several implications. First, while the chapter references problems in risk analysis and risk management that are generic and not specific to portfolio problems, it does not explore them in detail, e.g. processes for involving stakeholders, methods for assessing probabilities. Second, because it is about managing rather than just analyzing, decisions are relevant. Following Pate-Cornell and Dillon-Merrill [1], who describe the way that decision analysis overlaps with risk analysis and helps connect it to action, this chapter incorporates a decision analytic perspective, rather than a purely risk analytic perspective. Third, defining risk in terms of negative consequences, the chapter focuses on portfolio problems aimed at reducing risk rather than the many portfolio problems that merely involve risk (although the latter provide lessens applicable to the former). Finally, the range of portfolio and risk problems and methods found in practice is broad, and their intersection is not always well documented. Therefore, this chapter surveys broadly what the authors view as the relevant terrain and the general guidelines for navigating it.

Commonly, risk management is approached as efforts directed at single risks or ad hoc groupings of efforts directed at risks. Neither of these treatments explicitly considers possible interactions among projects or risks when risks stem from multiple sources. Whether a set of objects is considered explicitly as a portfolio is a choice made by whoever is considering them. Considering objects as a portfolio in this sense typically involves at least some additional effort in understanding relations involving multiple objects (risk, decision, etc.). Risks should be considered as a portfolio when risk from multiple sources must be managed in a coordinated manner. The need for coordination depends on the decisions under consideration: if the optimal decision that influences one risk could change depending on decisions or events that relate to another risk, then ideally those risks would be managed as portfolio.

Portfolios of risks arise in many domains, examples of which are summarized in Table 1. In real estate, a portfolio of properties faces the risk that one or more properties may lose value [2]. These losses may correlate with common initiating events such as local economic conditions or interest rates [3]. For business product portfolios, there is a risk that a new product will fail and the money invested in its development will be lost [4]. Of course, this is a risk that can be minimized by simply not launching any new products, and in this sense, it may be better viewed a risky portfolio of opportunities that involve risk rather than a portfolio of risks. In the oil and gas industry, there are portfolios of opportunities involving risk in drilling expensive, unsuccessful wells at one site that correlate with those of other sites [5]. Additionally, investments in information technology can be considered as a portfolio – more so than other corporate infrastructure investments since information technology shares both financial and non-financial resources in a coordinated manner [6]. Critical infrastructure within a region or of a certain type, such as dams, chemical plants, power plants, may be viewed as portfolios of assets with an associated portfolio of risks (due to failure) and associated investments that can also be viewed as portfolios. Purely financial assets (stocks and bonds) are commonly considered as portfolios with associated risks. Large scale engineering projects, such as those undertaken by NASA [7] may be viewed as portfolios of risks (involving multiple physical components and multiple program activities). Environmental risks may be regarded as a portfolio if shared hazards, pathways, or locations exist. In healthcare, a portfolio of diseases may result in a portfolio of negative health consequences (pain, death, economic costs), and can be mediated by a portfolio of treatments.

Table 1: Examples of portfolios of risks and risky portfolios

Problem \application	**Elements of portfolio directly affected by decisions**	**Nature of risk**	**Interactions between elements**
NASA unmanned flight	Design elements	Component failure leading to system failure	Causation, shared vulnerabilities
Oil and gas	Wells, seismic tests	Failure to find oil or gas leads to lost investment	Correlation in well performance, dynamic information
Infrastructure	Structures (repair and improvement)	Failures of any part of system lead to a variety of consequences	Shared vulnerabilities, causation
Business products	Individual products, project development phases	Loss of investment, disappointment	Production and market synergy or dissynergy
Real estate	Structures, land	Small and large scale economic, social and environmental conditions affect value	Virtuous and vicious cycles, supply/demand

The remainder of this chapter will focus on approaches to portfolio risk management and best practices associated with portfolio risk management. The following sections begin by reviewing key portfolio methods and how they can contribute to a quality decision process. Continuing in that line, next a comparison of the strengths and weaknesses of the different approaches together with a method of combining approaches is developed. The final section illustrates risk portfolio management concepts using the example of managing a portfolio of infrastructure risks.

2. APPROACHES TO PORTFOLIOS

It is useful to first consider four primary approaches to managing risk portfolios: project portfolio management, modern portfolio theory, risk analysis and risk management, and systems modeling.

These approaches have different strengths, and the manager who is aware of the range of possibilities can and do select among features of each of them (or other methods) as needed to apply an integrated approach to risk portfolio management. For this discussion, it helps to characterize the approaches in terms of the generic steps of a quality decision process. According to Howard [8], a decision process has six elements on which its quality depends: framing, alternatives, information, values, logic, and implementation. These will be highlighted as they arise in the discussion of each approach.

Framing seeks to define what is to be decided and why. This often means delineating the current decision from previously established policies and future downstream decisions. Within risk portfolio management, framing comprises defining which risks are to be included within the portfolio and clarifying any portfolio constraints. The generation of alternatives ultimately creates the decision problem. A good set of alternatives is creative but feasible; in the context of risk portfolio management, alternatives are the many means of reducing the risks (either the likelihood or the magnitude of negative consequences) within the frame. Information, ideally high quality information which is relevant, correct and complete, is necessary to make predictions. Risk portfolio management information typically pertains to individual risks. The values stakeholders and decision makers ascribe to the range of outcomes must then be determined. A logical synthesis of all of the above results in the identification and selection of the preferred alternatives. Finally, implementation of the selected alternative contributes to the realization of a desirable outcome.

The following rest of this section present sketches the four portfolio approaches in some more detail, together with some of their major variations. The goal here is to orient the reader but not to cover in depth all the issues that arise in practice, so many details are left out. Sections 3 and 4 consider these approaches as part of an integrated approach to risk portfolio management.

2.1 Project portfolio management

Project Portfolio Management (PPM) refers to an approach that selects which projects from a portfolio of candidates should be funded in the face of a budget constraint, often utilizing decision analysis methods. In the simplest case, there are no significant interactions between projects other than competition for funds. Projects are ranked in terms of a productivity index measuring their payoff per unit of resource used. When total cost is plotted against and total value for each portfolio, then the cumulative cost and value for the projects in order of productivity forms an efficient frontier for the set of possible portfolios. This approach has been widely applied in industry for R&D projects, especially in the pharmaceutical industry [9], and within the government for Army base closures [10]. In its simplest form, PPM finds the optimal portfolio solving $\text{Max}_{F_i}\ V = \Sigma V_i\ F_i$ s.t. $\Sigma F_i <= B$, where F_i denotes the funds allocated to project i, V_i denotes the value (or expected value) of project i and B denotes the available budget. At its most complex, PPM relies on sophisticated optimization tools to choose a strategy of coordinated and in some cases dynamic decisions subject to maximize some value measure on the distribution over possible outcomes subject to constraints describing resources, allowable interactions between actions, allowable levels on certain outcomes, laws of nature and many other properties of a situation. Decision support software may be useful in finding solutions.

PPM may be appropriate for managing risks when: the only (or main) interaction between risks is that they compete for resources; the objective is to minimize the relevant measure of risk, expected loss (compared to some baseline); there is a budget constraint for investing in risk reduction; for each risk there is a corresponding project to reduce that risk. Here, expected loss is the sum over all risks of the product of the probability that the negative outcome associated with the risk occurs and the magnitude of the outcome. The risks thus map to the project funding decisions. The alternatives in the portfolio decision problem are the levels at which to fund (if at all) efforts toward reducing each risk. PPM prioritizes the projects that reduce risk. This approach is used for business investments with some uncertainty, notably R&D. Typically, one project affects one risk. The valuation metric is typically expected dollar value or expected net present value; with portfolios of risks it may be natural to think instead of expected loss.

Consequences can take many forms, and it may be unnatural to consider only financial loss. Multi-criteria methods [11], [12] in PPM allow for project values (and expected values) across a range of objectives to be converted to a single score and expected value can still be used.

The portfolio is defined to include as its elements the possible activities under the funding authority so different levels can view the portfolio differently. For example, in the pharmaceutical industry, there may be a division focusing on heart-disease and the manager of this division views the portfolio as the products and compounds under development. At a higher level of the organization, an executive may think about a portfolio of business lines, e.g., cancer therapies, psychiatric drugs, medical devices, generic drugs, and medical supplies.

Framing for the problem involves determining the set elements that are candidates to receive funding. The primary information needed is the cost, probability of failure and the loss given failure (alternatively, the probability of success and the gain given success) associated with each project. If needed, interactions between projects are considered explicitly, e.g., two projects may be able share the cost of a common item, or two projects may be mutually exclusive. The decision rule that logically synthesizes all of these inputs is to select the portfolio that maximizes the measure of value subject to the budget constraint. If there are no interactions among risks other than their competition for funds, a good decision rule may be as simple as ranking projects in terms of a productivity index, e.g., value/cost, and funding them in order until the budget is exhausted.

2.1.1 Portfolio Risk, Variance and Non-Linear Loss Functions

One step in complexity beyond basic PPM is when outcomes for the elements of the portfolio interact in terms of a loss function, meaning the total loss depends on some measure of the sum of the realized consequences of the risks. If the combined loss from two risks combined is equal to the sum of the loss from each risk alone, then the *expected* loss of the two risks is equal to the sum of the expected loss for each risk alone. With a more general non-linear utility or loss function, this is not the case. One risk might be acceptable if there are no others, but if there is already a good possibility of ending up with a significant loss, the prospect of further loss may make render an additional risk unacceptable. In such cases, it's particularly important to understand the loss function/risk tolerance as there may be heuristics to manage risks based on it [13], [14]. For example, such a loss function can make risk diversification strategies attractive when they reduce the likelihood of extreme downsides. This feature complicates the portfolio management approach in two ways: a loss (or utility) function must be assessed, and computation of the expected loss minimizing portfolio has a non-linear element that may require more advanced algorithms.

One key difference between portfolios of risk reduction efforts and portfolios of, say, R&D projects is the way we would think about variance. For the former, the expected payoff of an investment to remove a risk is equal to its expected loss avoided. This is the probability of the negative consequence occurring in the absence of the investment multiplied by the impact of that negative consequence if it occurs. In a sense, the project is "successful" and adds value in the counterfactual situation where the negative would have occurred were it not for the investment, and otherwise it adds no value. In contrast, an R&D is successful and adds value if the technology works. If the project is not undertaken, there is no chance of it working. The variance for the value of a project in either case is the x^2 p(1-p) where p is the probability of the investment being a "success" and x is the "gain" in the case of a success.

With portfolio variance, there is a key difference. In the simple case of uncorrelated elements, the variance of the value of the portfolio is the sum of the variance of all its uncertain elements. R&D managers are concerned with the variance of the payoff of R&D, but with a risk portfolio, stakeholders are concerned with the variance of the results of the negative consequences. If no projects are funded, each risk contributes its variance to the total; the total variance of the risk portfolio is thus its original variance less the variance of the projects. If a risk-adjusted bang for the buck ranking is used, it may be necessary to add an adjustment factor for the projects in a risk portfolio, instead of subtracting it as in an R&D portfolio. In some cases, this implies that a high-variance portfolio of risk mitigation activities is desirable. This seeming paradox is due to a definition of value whereby the project value is negatively correlated with the value of the underlying portfolio of assets.

2.1.2 Interactions Within A Project Portfolio

Beyond interacting due to a non-linear loss function, risks can interact in more specific ways: risks may interact dynamically (whereby one risk leads to another)[1]. Risks may interact with respect to multiple criteria so that multi-attribute utility functions [15] might be suitable – especially if negative outcomes in one dimension make negative outcomes in another dimension more costly. There may be cost synergies or dissynergies or more general non-linear relationships between efforts and risk reduction. The logical synthesis of the project level information must produce a picture of the portfolio level risk. Advanced techniques in project portfolio management can accommodate some such considerations by expanding value calculations for specific projects, but at some point these interactions become too numerous and

[1] It may be possible in such cases to characterize and calculate outcome distributions based on these interactions using precedence matrices, e.g., [33].

complex to deal with on a project-to-project basis, and other methods are suitable for characterizing the situation and making decisions.

Interactions also present a major challenge for selecting actions within portfolio management (whether PPM and other approaches) because there can be a vast range of possible strategies involving a large number of decision variables. Combinatorial optimization (often using integer, binary and logic programming methods) and constrained optimization methods have been successful in this regard. Contingent portfolio programming [16] has been applied to risky projects with multi-stage decisions and with multiple objectives incorporating risk aversion. Most computational optimization methods require well-specified objective functions objectives can be identified and agreed upon. Robust portfolio modeling [17] can be useful where multiple objectives apply but where, as is common in portfolio problems, it is hard for the range of stakeholders to specify an exact objective function. There is a quite substantial literature on multi-criteria optimization methods [18], which portfolio managers.

The example at the end of this chapter illustrates different types of interactions between individual risks and projects. If there is rich enough interaction between them, enumerative approaches may be impractical. When there are well-defined correlations between many risks and where it is desirable to reduce the variance of the sum of the risk outcomes, financial theory described in the next section can be useful characterizing the portfolio.

2.2 Modern Portfolio Theory

Modern Portfolio Theory (MPT) was first formalized in several well-known articles [19], [20], [21], [22] and is now standard in finance textbooks, e.g., [23]. The crown jewels of MPT are elegant models resulting in simple formulas for investing under the assumption of efficient markets.

MPT is the nearly universal basis for management involving portfolios of risky financial assets, including stocks, bonds, and cash. Individual assets such as stocks have uncertain returns with both upside potential and downside risk. An investment portfolio consists of a number of different assets purchased. It is common to denote as w_i the weight of asset i in the portfolio. The return on an asset i, r_i, is the percentage by which its value changes over a period of time (e.g., annual return). Arithmetic shows that the return on a portfolio is the weighted sum of the returns of the assets in it. The *expected* return on a portfolio is likewise the weighted sum of the expected returns of its assets.

Due to the diminishing marginal utility of wealth (as well due to psychology and to the planning benefits of predictability), investors are generally risk averse. In market portfolios, this means that they would be willing to give up some expected return to reduce the spread of outcomes (typically measured as variance). Investors thus typically aim to achieve the maximum return for a given level of risk, or alternatively, to have the minimum risk associated with their expected level of return. For such investors, diversification is a free lunch. Variance can be reduced without affecting average return. If asset returns are uncorrelated, the portfolio variance is the sum of the product of each individual asset's variance and the square of the asset's weight in the portfolio. Then portfolio variance can be reduced arbitrarily close to zero by having many assets each with a weight in the portfolio of close to zero.

The fact that returns on assets are correlated makes this theory more interesting. Many investment risks are essentially uncorrelated, e.g., the fortunes of one company in one industry may depend on the occurrence an unpredictable technical breakthrough, while the fortunes of another may depend on the status of international relations. The insurance aspect of pooling poorly correlated risks is the motivation for having mutual funds. Individual mutual fund investors need not be concerned with the risk associated with specific companies. Stock prices, however, are often correlated. In a strong economy, for example, most stocks tend to go up while in a weak economy they don't. The correlations are even stronger between stocks of similar companies in similar industries. Variances and covariances of individual asset returns are estimated using a mix judgment and sophisticated econometric methods; unfortunately, MPT models require correlations and uncertainty looking forward to the future, but automated methods can only mine historic data. From these estimates, the variance of portfolio return can be calculated based on the weights in different assets and the matrix of covariances between the assets.

As a mathematical problem, it's theoretically possible to find the variance minimizing mix of investments for any level of expected return, i.e. the efficient risk-return frontier. With a substantial number of potential investments, however, finding this frontier could be a combinatorially large or even impractical problem. Fortunately, mathematical simplification is possible. The market as a whole is the sum of all assets weighted by their proportion of the market's total value that they represent. The Nobel prize earning references above show how useful this fact is. For purposes of constructing an efficient portfolio, it's often sufficient to summarize all the asset correlations in terms of the correlation between each asset's return and the total market return. The asset's correlation with the market as a whole is its β value (because the theory derives from statistical regression model coefficients), and the portfolio itself has a β value equal to the weighted sum of the β values of its elements. Thus correlation coefficients are a widely used and theoretically supported model simplification of multivariate probability distributions that would be difficult to assess directly.

This gives rise to the distinction between systematic and non-systematic (sometimes called idiosyncratic) risk. If the market is efficient (as is often assumed in theory), there is no way to hold stocks that give the same return as the market without also at least having the risk of the overall market. That risk, which cannot be diversified away, is called systematic risk. In contrast, a portfolio on the efficient frontier would have small proportions of many different stocks so that the idiosyncratic risk of each asset is diversified away.

But the investor can do better than holding a mix of risky assets. A diversified index of the market as a whole is on the efficient frontier. But the investor can (in theory) also hold a 'risk-free" investment – one with guaranteed return. U.S. Treasury bills have sometimes been assumed to be risk free in this sense. For any level of risk the investor might wish to accept, some percentage β of funds can be invested in the market portfolio and the remainder invested in the risk free asset leaves the portfolio to leave that risk. The variance of this investment mix is β times the variance of the market portfolio. In an efficient market, the expected return on this portfolio is $(1- \beta)\ r_f + \beta\ r_m$, where r_f is the rate of return on the risk free rate and r_m is the return on the market. The set of possible risk-return points obtainable by choosing different levels of the risk free investment is called the capital allocation line, and for each level of risk, the corresponding point on this line has return greater than or equal to the best portfolio that doesn't contain the risk-free asset. The value of β represents the leverage of the portfolio and could range from less than 0 (for a portfolio which actually moves in the opposite to the market, i.e., one in which stock is sold short in order to purchase the risk free investment) to more than 1 (a leveraged portfolio). The risk appetite of the total pool of investors determines just how high r_m must be in order for there to be an equal number of buyers and sellers of risk, and MPT says that current market prices will adjust so that expected r_m is at the right level.

Finding optimal financial portfolios is by now a well and continually studied problem. Simple algorithms often suffice, although professional investors such as hedge funds certainly have more complex and sophisticated models. These models may have constraints or value measures involving risk statistics other than standard deviation, e.g., value-at-risk, which gives the amount of loss associated with a given percentile. Financial engineering tools extend MPT concepts as well as options theory [24] to manage risk and return to sculpt a relatively desirable probability distribution by diversifying across assets, diluting risks, and creating or obtaining various insurance mechanisms defined in terms of future asset prices.

Arbitrage pricing theory, a fairly recent development, allows the portfolio to be further characterized in terms of multiple "β" factors, each corresponding to a different systematic and non-diversifiable risk (e.g., inflation risk, real estate risk and classic stock/bond market risk), with associated market equilibrium prices for each type of risk. This matches up nicely with applications of MCDA in finance [25], [26], where alternatives are evaluated in terms of a multicriteria model and the criteria correspond to sensitivity to different kinds of risks.

MPT-based approaches can be adapted to more general decisions involving risks. Smith and Nau [27] and others have shown how in decision analysis, some risks may be treated as market-based and assigned a market derived price while others risks must remain internal. Decision theorists [28] have identified utility functions with intuitively desirable properties that can also be characterized as simple functions of risk and return. In particular, Bell argues that $u(w) = w - be^{-cw}$ is an especially useful form, especially if risk takes the form of a normally distributed increment to initial wealth. Such a function based on mean and variance provides a bridge from the methods and statistics utilized by financial theory (as described above) to the concerns of loss-sensitive decision making. Because utility is usually decreasing in risk, the idea of efficient risk-return frontiers may also apply to non-financial portfolios of actions, and some aspects of the MPT conceptual framework have been used to describe such situations [29].

To put this in terms of the decision quality framework, in MPT the frame that defines what is to be included in the portfolio is primarily the notion of correlation, that is, potential investments are considered as parts of a portfolio to the extent that the investor is concerned about their correlation or lack thereof. Risks and decisions involve available investment instruments. Where the ultimate values may be correlated, where the level of correlation matters, and where there is some way to estimate the correlations, it is useful to consider these possible investments as the elements of the portfolio. The alternatives are best described in terms of how much of each risky entity to hold and in what overall mix. MPT typically values the portfolio solely in terms of risk (e.g., variance) and return (e.g., expected value, or expected rate of growth), and optimization formulas and algorithms logically synthesize the information with an aim to maximizing return for a given level of risk.

2.3 Risk Analysis and Risk Management

Risk analysis involves "detailed examination including risk assessment, risk evaluation, and risk management alternatives, performed to understand the nature of unwanted, negative consequences to human life, health, property, or the environment; an analytical process to provide information regarding undesirable events; the process of quantification of the probabilities and expected consequences for identified risks" [30], and **risk assessment** is "the process of establishing information regarding levels of a risk and/or levels of risk for an individual, group, society, or the environment." Risk assessment uses a set of tools to incorporate technical information and expert judgment about the structure, likelihood and magnitude of consequences from sets of related events. Ideally, it yields precise understanding of what drives risks and generates defensible probability distributions over outcomes from those risks under various scenarios. **Risk evaluation** (following the definition above) should aggregate information from the assessments associated with the portfolio of individual risks to develop a risk profile (e.g., a probability distribution) for the portfolio. It then relates the risk profile to the situation at hand (e.g., levels of background risk, importance of different potential losses associated with negative outcomes in various dimension). **Risk management** is the process of judging the desirability of potential consequences and selecting action alternatives, often in response to the findings of risk assessment and risk evaluation. It uses a loose set of tools, practices and processes, both qualitative and quantitative, that may be used with or without risk assessment. Risk management helps to ensure that there is a reasonable plan to reduce risks and that this plan is executed. The reader should note that the labels of these different activities are not always interpreted precisely and that there is overlap with other activity labels such as risk characterization and risk integration.

Risk assessment and risk management have been separated in many regulatory and planning applications. Risk assessment is considered to be a scientific evaluation process performed by scientists while risk management is associated with policy and regulatory process and is performed by regulators or planners. For example, the Environmental Protection Agency [31] specifically states that "risk assessment provides "INFORMATION" on potential health or ecological risks, and risk management is the "ACTION" taken based on consideration of that and other information" including, among others, economic, social and political factors as well as stakeholder values.

This artificial separation can result in practice in a disconnect between risk assessment and risk management. For example, the framing of problems in risk analysis usually is limited to identifying population and resources at risk and not risk management alternatives. It is generally done using conceptual models (influence diagrams visualizing relationship among stressors and resources at risk). The goal of risk assessment is not to prioritize risk management alternatives, but rather to assess risks resulting from multiple stressors and inform risk managers about these risks. Alternatives are not explicitly considered as a part of risk assessment; instead risk models are used to assess residual risks associated with implementing risk management alternatives that are developed outside of risk analysis process.

The focus for risk assessment is collecting information and developing models to assess risk associated with potential stressors. Risk assessments usually include steps of hazard Identification, exposure assessment, effects assessment and risk characterization. Hazards are usually identified utilizing historical data, reconnaissance site survey, preliminary models and other available information based on professional judgment of risk assessors involved in the process. In risk assessment conducted for regulatory compliance (e.g., Superfund sites in the US), this step also involves detailed consultations with stakeholders and regulators. Exposure assessment includes evaluation of the degree to which different population groups, environmental receptors and resources are affected by stressors while effects assessment quantifies the level of harm resulting from exposure to stressors. Risk characterization combines results of exposure assessments and effects assessments and compares resulting risk to regulatory benchmarks or other criteria deemed to be acceptable. Estimates and judgments related to uncertainty are part of the process, but these are rarely integrated explicitly and quantitatively during risk characterization.

Risk assessment informs risk management. Treatment of risk assessment process as “science” and risk management process as “policy” has resulted in very limited use of quantitative tools at risk management stage of risk analysis. Risk management alternatives are usually generated based on professional judgment and residual risks associated with their implementation is assessed using developed risk models. Values and other factors are incorporated in an ad-hoc manner and decision analysis tools for value assessment and for logical synthesis are rarely used. Alternative selection is extensively vetted through stakeholder consultations. Implementation of selected alternative usually requires monitoring to confirm expected reduction of risks. In many cases, though, expected risk reduction is not achievable and risk assessment needs to be repeated to meet regulatory needs [32].

In general, risk assessment treats individual risk independently and passes this information for consideration by risk managers. Sometimes simple integrated metrics are used (e.g., hazard index is developed as sum of hazard quotients for individual chemicals). In contrast to the simple portfolio problem, where the only decision variable available to management is allocating funds to projects, risk management has a richer set of alternatives about how to manage each of the risks, e.g., *mitigation*, *amelioration*, *diversification*, *risk sharing* (in which case, we may also be concerned with incentive effects, as in [33] as well as *establishing reserves* [34], *insuring,* or *obtaining information*, but the framework for quantitative assessment is not developed. In many cases single risk driver can be identified since risks, more so than project portfolios, tend to involve low-probability, large magnitude events; in some cases, we can essentially ignore their interaction because of this, while in other cases, their interaction is plausible enough that we have to worry about disastrous consequences of multiple large magnitude events that overload our capacity to deal with them. Assessment of low probability events and of correlations between those events thus becomes important. If we plot resources used against the measure of risk, we obtain a risk-reduction trajectory (similar to the Pareto-frontier for project portfolios).

2.4 Systems Modeling and Design

Interactions among a set of risks may be too complex and dynamic to practically capture with the approaches above. In that case, the portfolio of risks can be viewed as a system [35]. **Risk-based design** approaches are a more powerful tool for characterizing the problem structure. Dillon and Pate-Cornell describe a system for programmatic risk analysis where risks can be characterized by interactions among subsystems and components. Originally developed for NASA, their methodology divides optimization within a single project into four steps. The first identifies all viable alternatives, developing a minimum cost and residual budget for each. Next, a project specific residual budget curve is developed which allocates the residual budget between system reinforcement and project reserves using probabilistic risk assessment. Problem scenarios are then identified together with their likelihood to develop a managerial risk profile. Finally, the components are optimized together, with risk thresholds serving as constraints, to determine if an optimal project exists.

A variation on this approach accounts for project interactions within a project portfolio. Prospects for projects dependent on previous outcomes are calculated conditional on those previous outcomes to form an expected project cost. Initial projects incur a cost penalty, the difference between the non-conditional and conditional cost of the later projects, for disrupting later projects. Optimization occurs to minimize the second project expected cost and to minimize the cost penalty. The framework was originally developed for sequential projects directly affected by a project failure, such as an initial space orbiter mission followed by a lander mission [36]. However, projects occurring in parallel may also be considered. Parallel risks are represented in the projects' respective fault trees; a secondary tree can be developed to show the interaction among all projects. For a range of circumstances, assessments are performed to estimate the probability of the system being in a state of failure. The tree methodology treats each node as a random variable, with arrows showing probabilistic dependence. The final node shows the probability of project failure. This approach models the interaction among projects, information acquisition possibilities, and the whole range of scenarios affecting systems, with probabilistic risk assessment.

Risk-based design incorporates risk-analysis methods within the much broader discipline of **systems engineering** [37]. Systems engineering defines a system as a set of interdependent components (this set itself could be considered as a portfolio), aligned for a common objective or objectives. It is used to manage all aspects of complex projects and, therefore, could encompass an even broader notion of portfolio and of risk.

Events and consequences are quite often embedded within a system, e.g. a nuclear waste remediation system [38]. Thus, it is common to bring some systems modeling tools into any of the aforementioned approaches. In even more complex settings, **dynamic systems modeling** is useful, e.g. understanding feedback loops may be necessary in order to anticipate emergent events. For example, it may be that decision makers failed to anticipate the financial crisis that emerged in 2008 because they failed to take such dynamics into consideration. The global climate is an even larger system for which large dynamic models are used to gain insight about risks. The decision variables in such large systems may be a portfolio of interventions that modify feedback loops, etc., but these are not associated in any simple manner with specific initiating events, assets or consequences. Such modeling approaches go well beyond the problem of managing a portfolio of risks and are well-detailed elsewhere [39]. Overall, managers can choose from a wide range of approaches to craft an appropriate approach to managing risk portfolios with attention to the specific challenges of the situation they face.

3. STRENGTHS AND WEAKNESSES OF, CONNECTIONS BETWEEN, AND CHOOSING AMONG METHODS.

This section aids in selecting the right risk portfolio approach by comparing the relative strengths of the different methods and considering the type of situation in which those strengths are needed. It is helpful to start the discussion with Table 2, which summarizes (the authors' view on) the differences between the approaches described above as well as the characteristics of an integrated approach to risk portfolio management.

The different approaches have contrasting strengths and weaknesses. Project portfolio management identifies efficient allocation of resources, sometimes very much s, e.g., with dynamic models involving multiple resources. It tends to foster an organizational culture of efficiency too, and may explicitly encourage risk taking more than other approaches. On the other hand, although it is possible to account for project interactions, it's not simple to do so. The approach works simply when straight expected value is the decision criterion, but with risk portfolios, this may not be the case.

Table 2: Comparison of approaches

Process	Project portfolio management	Modern portfolio theory	Simple risk analysis and risk management	Systems Modeling	Integrated approaches to risk portfolio management
Framing **-Factors determin-ing what is treated as a portfolio** **-Constraining factors**	Funding authority and cycle Budget and other resources	Fungible / tradable market assets Ratios, funds available	Potential events and scenarios Acceptable risk levels, budget	Interactions between projects Acceptable risk levels, budget	Flexible Resources and performance requirements
Alternatives	Funding for each project	Percentage mix of assets	Tactics at project level	Allocating funds btwn projects/tasks	Look at many dimensions
Information and interactions	Required re-sources, prob-ability and value of suc-cess, ad hoc interactions	Return correlations and distributions, prices	Failure trees	Models of system and components	Decision-oriented interactions of various types
Values	Expected value or utility /MCDA	Risk and expected return	Acceptability based on probability of failure and magnitude of loss	Acceptability based on probability of failure and magnitude of loss	Multi-criteria / utility function selected for good risk properties
Logical synthesis	Optimization to maximize portfolio-wide productivity index	Compute theoretically optimal balanced portfolios	Simulation, scenario planning	Assign optimal distribution of funding to maximize objective	Compute as needed
Implementation	Project budgets and milestones	Trades and purchases	Processes, monitoring & responsibility	Budget development	One or more of other approaches

MPT methods are quantitatively tight so they can be used on a large scale where appropriate markets and mechanisms exist to allow for precise application of controls. However, as the current (2009) world-wide financial crisis demonstrates, it tends to miss system dynamics – and if conditions have changed, it may use the wrong assumptions (e.g., volatility is higher than assumed) and its decisions may not be robust in the face of such problems. Furthermore, it has not been adapted much to use beyond pure profit measures. Risk analysis and management deal especially well with uncertainty about events, as well as acknowledging the where there is ignorance. It adapts well to a system view and its processes (in theory) accommodate multiple stakeholders (as there commonly are in portfolio settings). On the other hand, risk management tools such as red-yellow-green classifications provide only a coarse prioritization. Simplistic risk management practices can be hard to translate to a specific implementation; they can lead to very (overly) conservative practices (e.g. where any probability at all of an event moves it out of the green category into yellow), and may lack sufficient quantification identify a most optimal plan. In order to tradeoff risks against benefits, MCDA approaches again can be useful.

In choosing the right tools, practitioners should match the strengths and weaknesses of different aspects of each approach to the needs of the situation. For example, certain established approaches combine two or more methods. In terms of risk measures, *value-at-risk* [40] combines from risk analysis and finance to summarize and communicate risk making it suitable for managing certain day-to-day operations, although as a management tool, it does not capture well problems involving system failure and very low probability events. *Risk-based design* combines risk analysis and systems analysis / design, and it is useful for capturing complex interactions between risks and providing a productive approach to managing risks. Programmatic risk analysis [41] combines elements of decision analysis and risk analysis in a customized approach to managing related risks. Decision analysis and risk management techniques can be combined to consider alternatives involving mitigation, amelioration, allocation and acceptance of risks. Finally, MCDA has been combined with risk analysis [42] to translate the range of consequences into a basis for decisions.

In an integrated approach, portfolio problems should be classified according to which tools are appropriate and then those tools should be used. The example in the following section illustrates this idea by describing a portfolio of risks (dams and levees) and discussing how they might be managed with an integrated approach and the other approaches explained above.

4. RISK PORTFOLIO MANAGEMENT: APPLICATION EXAMPLE.

This section uses an example from the US Army Corps of Engineers (USACE) in order to illustrate in a specific context the general steps for managing a risk portfolio and to discuss how, within in integrated approach, different methods can be fitted to a particular problem. We note that there is no single representative portfolio in terms of the specific methods that should be used to analyze and manage it, but the idea of actively selecting methods for managing each of the various dimensions of the problem is common across portfolios. USACE and related government agencies manage thousands of dams and levees. Within the dam and levee portfolio, these agencies are responsible for allocating limited budgets to reduce the risk of failure. However, if risks between projects correlate, or if losses are not constant functions, traditional methods of allocating funding may not maximize risk reductions. Portfolio methods must then be utilized to ensure the greatest possible risk reduction in these complex circumstances. Differing elements of the previously outlined risk approaches apply in the context of dams and levees.

The methods used to support decisions about the portfolio should match the degree and type of complexities in the risks, decisions and consequences of the portfolio and their interactions, e.g. correlations among risks and non-linear loss functions. Risk correlations occur when the likelihood and the impact of failures for different structures are not independent. Put similarly, the loss due to two related failures may cause much greater damage than the sum of the damage caused if both failed separately. For example, a ten-foot rise in water level may be more than twice as costly a five-foot rise or the failure of an upstream dam may precipitate the failure of a downstream dam. Correlations thus account for the tendency of issues to change with the surrounding events. Failures may be correlated with internal or external stressors. Structures with a common design or other common characteristics could fail disproportionately in response to a specific stressor. In particular, climate change could increase the likelihood that a stressor will affect a large number of structures (sea level rise may affect structures in low-lying areas and changing precipitation patterns could have similar effects on structures in similar locations). Therefore, considerations such as diversification, system design, and dynamics could be important and could incorporate qualitative or quantitative measures of performance, or in an integrated approach [43], [44]. Correlations add complexity to the management of projects by instilling relationships with surrounding projects and events. Non-constant loss functions also add complexity, as it is less intuitive to convert losses to an equivalent single dimensional value, such as monetary cost.

Managing dam and levee risks can be accomplished through both infrastructure improvements and policies/practices. Infrastructure improvements are projects that require substantial amounts of funding to put into place new hardware. These address risks to many types of structures and facilities, various groups of people, and risks from chains of events. Infrastructure improvements also ameliorate consequences ranging from loss of life to economic loss and social displacement to environmental damage. Within dam and levee management, there is a substantial but limited budget for infrastructure improvements.

Limited budgets for infrastructure improvements stimulate interest in other ways of achieving risk reduction with less financial expenditure, typically through policies and practices. Policies and practices involve zoning, organizations, and planning for both standard operations and contingencies. This complicated problem is surely better viewed as a portfolio than as a set of independent risks. What is done with regard to one risk will can influence other risks. For example, contingency planning to open flood gates during high water events could cause downstream flooding, having the effect of exchanging dam failure risk (probability) for human risk (level of harm). A combination of policy methods and infrastructure improvements will typically be involved in most risk reduction strategies.

Several courses of action are available to the USACE given the before mentioned considerations. The worst practice might be to fix each dam after it fails. Alternatively, a plan might be announced to perform all proposed repairs on all structures, but that plan would be unrealistic because it would likely exceed any available budget and some work would remain unfinished. Historically, engineers for each dam would write a justification for funding agencies about why they need money and decisions about allocations would then be made through a rather unstructured political process. But if it is taken as given that the societal objective ought to be to minimize damage, that approach also has great shortcomings. USACE has recently [45] taken steps toward more systematic decisions about these risks. Dams are graded on engineering criteria (e.g., cracks, voids, spilling, concrete deterioration, and impact criteria (i.e., magnitude of consequences – potential hectares flooded, people killed, etc.) as in [46]. The product of these two scores produces a prioritized list of dams (high, medium, low risk). Investment is made in the highest risk structures first. This approach imposes discipline on the process, but it may not be sufficient to allocate the available budget so as to minimize society's expected loss. It doesn't explicitly consider cost, so optimization algorithms cannot be applied.

Project portfolio management methods that value risk mitigation efforts can improve the effectiveness of resource allocation beyond simpler approaches that prioritize risks in terms of probability of failure or severity of failure. Where losses may be correlated, notions from financial portfolio management, MPT, can be applied. With coastal projects levees, there is reason to expect losses to be correlated along various measures and returns on investments to be correlated with respect to those measures. With dams, losses should correlate with rainfall and upstream dam failures. Notions from MPT have been applied these problems [29]. Small changes are more easily accommodated than large shifts, so there's reason for risk aversion in this kind of portfolio. This would be true – to different extents – with respect to shipping capacity, human displacement, ecological damage, etc. With the utilization of a multi-attribute utility function (ideally of a form that can be expressed in terms of risk and return) one can describe portfolios with respect to each of the measures.

The purpose of the following case discussion is to illustrate how the generic steps might translate in a realistic situation, as well as to illustrate how a computational model can be incorporated. The computational model used here is a relatively simple spreadsheet. For more complex applications, a more complex spreadsheet, a spreadsheet with Monte Carlo simulation [47] and optimization tools included in an add-in (e.g., Risk Solver by Frontline Systems, Decision Tools Suite by Palisade, or Crystal Ball Decision Optimizer by Oracle), or dedicated modeling software could be suitable.

The authors' list of generic steps (contributing to various elements of decision quality):

- Define what is in the portfolio and defining constraints (Framing)
- Define consequences of concern (Setting up to obtain information that informs values)
- Identify initiating events and emergent risks (Information)
- Structure and characterizing individual risks (Information)
- Map risks to decisions about steps to reduce risk (Alternatives)
- Characterize impact of steps to reduce risks (Information bridging from Alternatives)

- Characterize cost/resources of steps (Information bridging from Alternatives)
- Characterize interactions between risks (Information bridging to Logical Synthesis)
- Calculate portfolio risk profile (Logical Synthesis of Alternatives and Information)
- Calculate portfolio risk metrics / characterizing portfolio risk (Value)
- Choose a plan (Logical Synthesis of Frame, Alternatives, Information and Value)

In selecting the exact approach, it is useful to apply an underlying concept of decision quality. In each of the six dimensions, a 100% quality is defined as the point at which additional improvement stops being worth the additional effort required to achieve it. What this point is depends on the situation. So, while no dimension should be forgotten, it is possible to go beyond this point by doing too much analysis on some aspect of the decision. Thus, modeling structures and tools should be used that support the level detail appropriate (generally, this is still determined subjectively) for the problem in each dimension.

The following example, using a simulated portfolio (using random number formulas in a spreadsheet to generate the parameters for each structure's condition, size, cost of repair, potential impact of failure, etc.) demonstrates how the generic steps could apply in case of USACE's management of dam and levee risks.

Define what is in the portfolio
For this example, the portfolio set will be the dams and levees on the east coast. This selection is based partly on the charge and expertise of the USACE, but also on commonalities, concerns, consequences, and budgets shared across the projects. Also, it is unnecessary to expand the portfolio beyond this. If the goal is to manage and mitigate risk at dams and levees, including sewage processing systems or railways will not contribute to this goal. Though a greater portfolio would allow for greater resource sharing and more subtle analyses of risk, it will still contribute little to the ultimate goal of dam and levee risk reduction.

Define consequences of concern
If a dam or levee fails in part or completely (in reality, there are multiple failure modes some of which have more severe consequences than others), the immediate consequence is water release. That itself isn't a problem, though downstream that water could destroy cities, homes, and habitat. The main consequences can be classified as economic, social, or environmental. Any of the structures failing would have roughly similar consequences, although they would differ in magnitude. For other portfolios, it's possible only one description, financial, might be the only concern.

Identify initiating events and emergent risks
Negative consequences for dams and levees could be caused either by structural failure or operation out of specification. Operation out of specifications might result from stressed conditions such as changes in weather patterns (e.g. rainstorms). Internal causes could result from design flaws, maintenance problems, and age. In this portfolio, similar fault trees are assumed to simplify the analysis.

The logic in this simple model is that initiating events may occur (not shown in model), actions may be taken to reduce risks, outcomes for individual risks then result, and the sum of all risk outcomes determines the portfolio loss.

FIGURE 1: Drivers of portfolio value

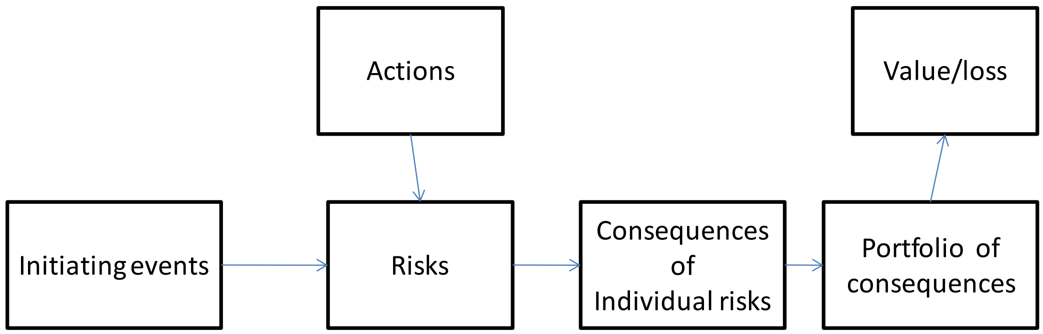

Structuring and characterizing individual risks

In the context of the example, there must be a determination about how to judge the severity of risk. Various approaches are applicable. A simple approach might be to label dams as low, medium, or high priority and levees as structurally sound, in need of repairs, or deficient. A second approach might be to prioritize structures based on their size or age. Yet another approach would be to classify dams and levees by the size of their downstream cities. For a PPM approach, if intuitive prioritization is not a fine enough process one can be more precise and specify a probability of failure to capture both the condition of and the stress on a given dam as wells as a numerical value for the consequence of failure, based on either a single value or multiple values such as number of displaced people and the business assets lost. In other cases, detailed fault trees might be useful. In either case, if failure probabilities are assigned, they must be defined with respect to a time period. Regardless of the method, with portfolios, it is desirable to be an honest broker by using a consistent method for all elements of the portfolio.

In Figure 3, below, the fourth column gives the probability of failure of each dam, while the magnitude of failure consequences is shown in the sixth column.

Map risks to decisions about steps to reduce risk

To manage the portfolio of risks, it's necessary to map the set of risks to a set of decisions. For the example, the simplest mapping method, assuming that there is a single project that will completely eliminate risks for a given dam or levee, was utilized. A more complicated and realistic mapping might include multiple decisions such as diverting water in addition to repairing the structure. In these more complicated situations, the model would require entries for the different decisions and their incremental changes to probabilities and magnitudes. This could mean a much longer list of decisions with logical constraints on the decisions and mapping the interactions among decisions.

In Figure 2, it is assumed that each project corresponds to a single reduction decision. As such, a full listing of decisions is not developed.

FIGURE 2 Logic of spreadsheet model

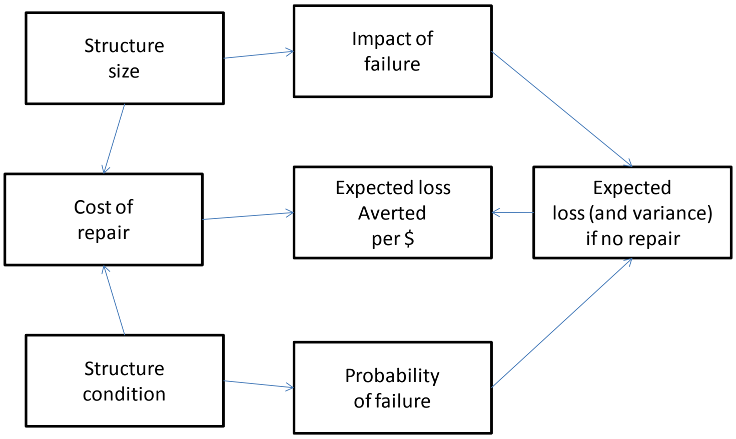

Characterize impact of steps to reduce risks

Steps to reduce risks can decrease the overall probability of failure, the probability of failure given an initiating event, or the consequences of failure. The anticipated decrease can then be compared against the no action case. The decrease in expected loss for each risk step is calculated as the difference between the original expected loss and the expected loss after the step is taken. In Table 3, it's assumed each step reduced the probability of failure to zero. A wider selection of possible risk reduction steps can also be accommodated. In a spreadsheet, this requires several extra columns and some additional calculations the find the impact of each step and the overall risk for any portfolio plan.

Table 3: A Simulated portfolio of risks in a spreadsheet model

Structure index	Structure size	Problem level	Prob of Failure	Cost to repair	Impact of failure	Expected loss	Variance of loss	Exp. loss averted per $	Rank	Funded?
1	3.27	4	0.13	0.56	3.57	0.47	1.46	0.04	18	0
2	10.35	3	0.49	2.15	19.72	9.65	97.18	3.95	1	1
3	6.98	2	0.15	2.96	13.21	2.01	22.55	0.29	9	0
4	4.32	2	0.31	1.39	4.28	1.31	3.89	0.24	10	0
5	8.30	2	0.41	0.85	8.91	3.69	19.25	1.37	3	1
6	8.78	3	0.32	2.90	6.22	1.98	8.39	0.64	4	1
7	2.26	4	0.11	0.57	2.53	0.28	0.62	0.09	15	0
8	9.96	3	0.40	3.81	3.57	1.41	3.05	0.23	11	0
9	8.16	2	0.05	0.94	28.96	1.48	40.64	0.42	6	1
10	1.71	3	0.25	0.39	1.17	0.29	0.25	0.07	17	0
11	10.04	1	0.15	0.91	6.06	0.91	4.71	0.41	7	1
12	5.31	3	0.46	1.15	1.21	0.56	0.36	0.11	14	0
13	1.87	4	0.00	0.30	1.98	0.01	0.01	0.00	22	0
14	3.64	2	0.04	1.07	1.24	0.05	0.06	0.01	20	0
15	1.97	2	0.10	0.54	1.20	0.13	0.13	0.04	19	0
16	5.65	2	0.27	0.62	2.23	0.61	0.98	0.20	12	0
17	2.32	3	0.03	0.60	1.36	0.04	0.06	0.01	21	0
18	7.52	3	0.36	1.21	14.89	5.43	51.37	1.56	2	1
19	5.29	3	0.00	0.70	1.15	0.00	0.00	0.00	24	0
20	6.66	4	0.01	0.87	23.95	0.20	4.66	0.09	16	0
21	4.39	2	0.00	0.44	3.05	0.00	0.01	0.00	23	0
22	4.54	4	0.38	0.98	9.14	3.47	19.69	0.56	5	1
23	5.07	1	0.28	1.24	7.09	1.97	10.08	0.39	8	0
24	8.87	4	0.00	1.00	61.93	0.00	0.11	0.00	25	0
25	11.08	2	0.17	2.83	3.13	0.52	1.36	0.18	13	0

Characterize cost/resources of steps

Effectively managing portfolios of risks involves allocating resources effectively. For dams and levees in the example, the repair budget is the only resource under consideration; estimating the cost associated with a given repair for a given structure is a standard engineering task. In other cases, there may be multiple resources that are limited and could not easily be exchanged for one another, e.g., in a public health emergency, there may be a fixed amount of hospital beds, medical personnel, and supplies and no way to get more of one resource by giving up (or even selling) another. The amount required of each of these resources by a risk reduction step should then be estimated.

Characterize interactions between risks

Characterizing risk interactions is a key and complex step. There may be interactions in terms of the likelihood of events (e.g., causation), the consequences of events, or the effectiveness of actions taken reduce either likelihood or magnitude of consequences of two or more risks.

In the basic numerical example, it's assumed for simplicity that there are no interactions between the risks. For dams and levees, this could be true if each structure were unique and located in an entirely different geographic region. In more complex situations, risks can interact in multiple ways. Examples include:

- If one structure fails, structures downstream are at greater risk of failure. Sometimes this can be accommodated in a spreadsheet using a precedence matrix.
- If one structure fails, it could mean that other structures with the same design are also more likely to fail. Such risk could be incorporated into a spreadsheet using joint probability or correlation tables.
- If there are dams on two rivers that join into a single river downstream, the damage caused by two dam failures together could be worse than the sum of damage each one failing alone. The two dams could be modeled as a single system with multiple failure modes. The interactions of parallel or serial elements of such systems are simply modeled using logic functions.
- Some of the initiating events that contribute to failure (e.g. particularly heavy rainfall over a wide region) will affect multiple structures. This could be incorporated into the model by assigning probabilities to different initiating events and then entering probability of failure conditional on the set of initiating events. The total probability of failure would end up being the result of a calculation rather than a directly entered value.
- If the consequences of failure are uncertain, consequences could be correlated across risks. For example, a general increase in population and urban development could increase the potential cost of flooding everywhere. Such consequences could be correlated within the model.

- There are synergies in risk reduction actions, for example, if a dam and a nearby levee are fixed, it might be cheaper to do them both together than the sum of their costs if considered independently, because construction equipment won't have to be transported.

Calculating the portfolio risk profile

Once all the individual risks, their interactions, and the steps to reduce the risks are characterized, it's meaningful to look at the overall range of outcomes. A typical format for the risk profile is a probability distribution over loss. If loss is multidimensional, the risk profile could take the form of a joint probability distribution. These probability distributions can be calculated, though in general they are often simulated. In some cases, statistics (especially mean and variance, as in financial portfolios) for the distribution suffice. For example, if the family of distribution is known or if the expected loss can be calculated directly from the statistics (e.g. with a linear loss function, expected portfolio loss is by definition the mean of the distribution). These statistics can often be calculated analytically, saving the effort of simulation. Whether this works depends on the risk measures and loss functions that apply in the following step. Profiles may also be intertemporal, capturing risk at multiple points in time [48], [49]. In the example, the expected loss across the whole portfolio, and thus the decrease in expected loss associated with a portfolio of risk reduction activities, is simply the sum of the expected losses associated with each risk. If projects are independent, the variance of the total portfolio loss, which can be calculated once assuming no reduction efforts and again with a portfolio of reduction efforts, is similarly equal to the sum of the variance of the loss for each risk i, $p_i(1-p_i)\ loss_i^2$. In cases with moderately complex correlations, variance for portfolios with correlated elements can also be calculated formulaically.

Calculating portfolio risk metrics

From the risk profile, one can calculate summary statistics – the sufficient statistics above and possibly other quantities to be calculated from the profile that will feed into evaluation of plans.

For the dam and levee example, economic, societal, and environmental costs are all of concern. A total (multi-criteria) damage measure could be calculated as a weighted sum of, say, hectares of habitat destroyed, number of people dislocated, and dollars of infrastructure lost. There are various methods in risk and decision analysis for estimating such tradeoff weights. The marginal harm to the area would become steeper as the magnitude of loss increases, as it becomes harder for regions to recover, so some non-linear utility (or loss) function is appropriate.

For illustrative purposes, the numerical example here uses only a single damage measure, and a simple risk-averse utility function (for illustrative purposes only) over just one metric in terms of mean and variance where utility u = 100 – expected loss – variance/10. Thus, if all risk were eliminated, utility would be 100, and it is potentially unbounded on the downside.

In some situations similar to this one, there are specific guidelines to plan with regard to the worst case that has a 1/100 chance of occurring. That would be similar to a value-at-risk measure, which has been used in financial planning. The “worst case” would be the first percentile of simulation output. Alternatively, there may be some limit above which the loss becomes catastrophic, for example so much habitat disappears that coastal species go extinct. In that case, the probability that the total habitat loss exceeds some threshold amount could also be a simulation output.

It is reasonable to apply multiple criteria to the basic spreadsheet using a simple formula so that the total loss for the portfolio is a weighted sum of the loss function for economic, social and environmental factors. More sophisticated multi-attribute utility functions are still relatively simple to compute with formulas that take attribute measures as inputs but using them requires careful encoding of value judgments.

An appropriate utility function for decision making uncertainty must reflect the decision maker’s / stakeholder’s preferences as follows: the expected utility for one risk profile is higher than that for another if and only if the decision maker prefers the first risk profile to the second. In risk situations, there may be no single clear decision maker and hence no single utility function can be entirely appropriate [50]. In that case, simpler and/or more robust measures could be used such as outranking-based multi-criteria decision analysis methods, or analysis using a range of utility function parameters could be presented. Recommendations may not vary much over the range of possible weightings. If they do, these results can still be informative to some further political process needed to ultimately choose priorities.

Choosing a risk management plan

Once the portfolio level measure can be calculated for any plan, the optimal plan must be selected. The gold-standard approach is mathematical programming or optimization about which there are many textbooks (e.g., [51]) to find the values of decision variables that maximize or minimize an objective function subject to a set of constraints. For risk portfolios, the objective could be: to minimize the expected value of the loss function or to minimize the probability that loss will exceed some level. The objective could involve one dimension or multiple dimensions. Decision variables typically allocate resources to activities or investments, but could also include scheduling and timing of efforts. Constraints typically include the available resources, but can also include policy constraints, e.g., the probability of losses exceeding x cannot be greater than y%. Constrained and combinatorial optimization methods range from relatively simple to highly sophisticated, requiring extensive data collection for inputs and extensive efforts at formulation. If optimization proves difficult, the choices are to simplify the model or to bring in a professional with special expertise in optimization methods.

For the simple dam case as developed here that ignores interactions between risks, optimization is relatively straightforward and doesn’t require search algorithms. In the basic case where the objective is simple expected loss and the constraint is a monetary budget, it works to rank projects in terms of a productivity index: expected loss avoided per dollar of investment and to select the top ranking projects up to the point that the budget is expended. The same method works if the loss associated with a failure is defined as the weighted sum of losses in each of the three dimensions, in which case the expected value of that weighted sum is used as the numerator of the productivity index[2]. If the utility function uses expected loss and variance only, a heuristic that modifies the investment's productivity to account for its impact on variance may suffice. Otherwise, a fairly simple binary search method will find the optimal portfolio.

Table 4 below shows the expected loss, variance and the other calculated measures under different prioritization rules. The columns labeled "do nothing" and "arbitrary choice" are included to give a baseline of what happens when there is no strategy, for the purpose of comparing the impact of the other strategies. Funding choices change with different rules, and as rules better approximate the maximization of portfolio utility, funding choices tend to converge toward what a gold-standard decision analysis would recommend. It is less helpful to rank choices on the basis of benefit-related characteristics that are correlated with cost in predictable ways. For example, structure size and cost are correlated, so there is no benefit to choosing either large structures – whose repair cost tends to be high, or low-cost repairs – which tend to impact only smaller structures.

Ranking projects in terms of the impact of failure or the probability of failure leads to a substantially better portfolio; ranking in terms of expected loss combines both of these factors and leads to still more improvement, and best of all by a significant margin is using the productivity index, expected loss averted per unit of cost. For all of these better rankings, there is not a great difference in which projects are funded. Ranking with multiple criteria is not shown here, but the implementation and implications are similar. If projects share initiating events, correlated responses, or causal interdependencies, the projects would start with a similar model with additional interactions structured to calculate the consequences for the set of individual risks. Simulation results could then be used to estimate the distribution of outcomes by which potential portfolios are valued. Note, correlations would typically be entered in the simulation tool itself, e.g., in specifying distributions for a cell and if necessary specifying correlations between various cells. It's possible to do this within a spreadsheet, e.g., by using a common random number as input to formulas for different cells, but such an approach is often unwieldy.

[2] Although appealing in its simplicity, the use of productivity indices can have substantial limitations. A single productivity number can more easily reflect cost efficiency than cost effectiveness, and may be insufficiently sensitive to context in cases where effectiveness differs from efficiency. Examples include dealing with uncertain costs, multiple resource constraints, multi-stage decision processes, synergies or dissynergies across expenditures, large discrete expenditures that must be coordinated within a fixed budget, or non-fixed budgets.

Table 4: Portfolio results under different prioritization rules

Criterion for funding	Do nothing	Arbitrary choice	Structure size	Problem level	Prob. of Failure	Cost to repair	Impact of failure	Expected loss	Exp. loss averted per $
Cumulative cost	0.00	9.86	9.70	9.77	8.94	9.26	9.13	8.14	9.94
Cumulative value	0.00	11.64	12.50	10.06	18.78	11.57	18.77	24.26	26.61
Expected loss	36.47	24.83	23.97	26.41	17.68	24.89	17.70	12.21	9.85
Final variance	290.89	203.55	184.59	210.77	151.35	198.40	74.38	80.84	49.65
Exp. Loss averted per $	0.00	1.18	1.29	1.03	2.10	1.25	2.06	2.98	2.68
Utility = 100-E(Loss)-Var/10	34.44	54.81	57.57	52.51	67.18	55.27	74.86	79.70	85.18
Which projects are funded?									
1	0	0	0	0	0	1	0	0	0
2	0	0	1	0	1	0	1	1	1
3	0	0	0	0	0	0	1	1	0
4	0	0	0	1	0	0	0	0	0
5	0	0	0	1	1	1	0	1	1
6	0	0	0	0	0	0	0	0	1
7	0	0	0	0	0	1	0	0	0
8	0	0	1	0	1	0	0	0	0
9	0	0	0	1	0	1	1	0	1
10	0	0	0	0	0	1	0	0	0
11	0	0	1	1	0	1	0	0	1
12	0	0	0	0	1	0	0	0	0
13	0	0	0	0	0	1	0	0	0
14	0	0	0	1	0	0	0	0	0
15	0	0	0	1	0	1	0	0	0
16	0	0	0	0	0	1	0	0	0
17	0	1	0	0	0	1	0	0	0
18	0	1	0	0	0	0	1	1	1
19	0	1	0	0	0	1	0	0	0
20	0	1	0	0	0	1	1	0	0
21	0	1	0	0	0	1	0	0	0
22	0	1	0	0	1	1	0	1	1
23	0	1	0	1	0	0	0	0	0
24	0	1	0	0	0	0	1	0	0
25	0	1	1	1	0	0	0	0	0

5. Summary

This chapter has described what makes a set of risks worth considering as a portfolio and how decisions about managing that portfolio of risks can be structured and executed. Some sort of quantitative modeling is almost always needed. Significant gains come just from recognizing portfolio management as an optimization problem, i.e., minimizing the expected value of the loss function by choice of risk management decisions subject to a resource constraint. Depending on the characteristics of the portfolio of risks and the potential actions to reduce them, different modeling approaches can be used. Simpler approaches have the benefit of taking far less time to perform and explain, but if they are too simple, they ignore important interactions between risks. This can lead to underestimating the remaining portfolio risk or overlooking ways to eliminate more risk with a fixed budget, or otherwise getting the wrong answer. The specific theoretical approaches described in this chapter are well-documented elsewhere in risk analysis, finance, and operations research. The reader should use this chapter as a guideline for identifying an approach and the steps to plan for it, and to go to the referenced literature for in-depth assessment and modeling guidelines.

This chapter has focused on assessment and modeling approaches for managing risk portfolios. There are certainly other related issues to keep in mind. To implement decisions, for example, requires a variety of managerial and relationship skills. It is surely helpful to have transparency in the process so that stakeholders understand the process and believe it to be fair and valid. The process itself may involve oversight, quality controls, and participation of key experts, decision makers. It can be helpful to question the robustness of models. This means asking what happens to the decisions if assumptions about key relationships or parameter values are wrong. Sensitivity analysis techniques are a good starting point for this. Such steps can both protect against oversights and errors in analysis, and create some of the trust required for implementation. Finally, implementation of efforts to reduce risk often involves handing off responsibility to engineering and project management professionals. Careful definition and documentation of exactly which alternatives affect which risks, to what degrees, at which times and with which resources will facilitate a smooth handoff.

Although explicitly managing a set of risks as a portfolio is more complicated than managing them merely as a series of individual risks, it is often still worth the additional effort.. In such diverse domains as financial management, product development, and military procurement, portfolio management tools have been explicitly adopted even to the point where there is a high level job with the title “portfolio manager.” Risk managers have long been managing implicit portfolios of risks, and they can productively incorporate a variety of portfolio management techniques.

Acknowledgements:

We would like to thank Alex Tkachuk, Laure Canis and Dr. Lambert for advice and discussions. This effort was sponsored in part by the Civil Works Basic Research Program by the U.S. Army Engineer Research and Development Center. Permission was granted by the USACE Chief of Engineers to publish this material.

References

[1] M.E. Pate-Cornell, and R.L. Dillon. The respective roles of risk and decision analyses in decision support, Decision Analysis 3:4 (2006), 1-13

[2] S. Lee, and S. Stevenson, Real estate portfolio construction and estimation risk, Journal of Property Investment & Finance 23:3 (2005), 234 – 253.

[3] S.C. Wooldridge, Balancing capital and condition: An emerging approach to facility investment strategy, Doctoral Thesis, Massachusetts Institute of Technology, 2002.

[4] M. Menke, Managing R&D for competitive advantage, Research Technology Management 40:6 (1997), 40-42.

[5] M. Skaf, Portfolio management in an upstream oil & gas organization, Interfaces 29:6 (1999), 84-104.

[6] R.J. Kauffman, and R. Sougstad, Risk management of contract portfolios in IT services: The profit-at-risk approach, Journal of Management Information Systems 25:1 (2008) 17 – 48.

[7] B. L. Botkin, Applying Financial Portfolio Analysis to Government Program Portfolios, Masters thesis, Naval Post-Graduate School, June, 2007.

[8] R.A. Howard, Decision analysis: Practice and promise, Management Science 34:6 (1988), 679-695.

[9] P.T. Sharpe, and T. Keelin, How SmithKline Beecham makes better resource-allocation decisions, Harvard Business Review 76:2 (1998), 45-57.

[10] P.L. Ewing, W. Tarantino, G.S. Parnell, Use of decision analysis in the army base realignment and closure (BRAC) 2005 military value analysis, Decision Analysis 3:1 (2006), 33–49.

[11] V. Belton, T. Stewart, Multiple criteria decision analysis: An Integrated Approach, Springer, 2001

[12] M. Zeleny. Multiple criteria decision making. McGraw-Hill, New York, 1982.

[13] P. Delquié, Interpretation of the risk tolerance coefficient in terms of maximum acceptable loss, Decision Analysis 5:1 (2008) 5–9.

[14] J.L. Ringuest, S.B. Graves and R.H. Case, Mean–Gini analysis in R&D portfolio selection, European Journal of Operational Research 154:1 (2004), 157-169.

[15] R.L. Keeney, and H. Raiffa, Decisions with multiple objectives: preferences and value tradeoffs. Wiley, New York, 1976.

[16] J. Gustafsson, A. Salo, Contingent Portfolio Programming for the Management of Risky Projects, Operations Research 53:6 (2005), 946-956.

[17] J. Liesiö, P. Mild, A. Salo, Preference programming for robust portfolio modeling and project selection, European Journal of Operational Research 181:3 (2007), 1488-1505.

[18] M. Ehrgott, X. Gandibleux, Multiple criteria optimization: State of the art annotated bibliographic surveys. Kluwer, 2002.

[19] H. M. Markowitz, Portfolio selection, Journal of Finance 7:1 (1952), 77–91.

[20] W.F. Sharpe, Capital asset prices: A theory of market equilibrium under conditions of risk, Journal of Finance 19:3 (1964), 425-442.

[21] W.F. Sharpe, A simplified model for portfolio analysis, Management Science 9:2 (1963), 277-293.

[22] F. Modigliani, and M. H. Miller, The cost of capital, corporation finance and the theory of investment, American Economic Review 48: 3 (1958), 261-297.

[23] A. Brealey, A., and S. C. Myers, Principles of Corporate Finance, 7th edition. McGraw Hill, New York, 2003.

[24] F. Black, and M. Scholes, The pricing of options and corporate liabilities, Journal of Political Economy 81:3 (1973), 637–654.

[25] W. Hallerbach, and J. Spronk, A multi-dimensional framework for financial-economic decisions. Journal of Multi-Criteria Decision Analysis 11:4-5 (2002), 111-124.

[26] W. Hallerbach, and J. Spronk, The relevance of MCDM for financial decisions, Journal of Multi-Criteria Decision Analysis 11:4-5 (2002), 187-195.

[27] J.E. Smith, and R.F. Nau, Valuing risky projects: Option pricing theory and decision analysis, Management Science 41:5 (1995), 795-816.

[28] D. E. Bell, Risk, return, and utility. Management Science 41:1 (1995), 23-30.

[29] J.C. J. H. Aerts, W. Botzen, A. van der Veen, J. Krywkow, and S. Werners, Dealing with uncertainty in flood management through diversification. Ecology and Society 13:1 (2008), Article 41. URL: http://www.ecologyandsociety.org/vol13/iss1/art41/ accessed 08/05/09

[30] Society for Risk Analysis, online accessed 8/2/09, http://sra.org/resources_glossary_p-r.php.

[31] Environmental Protection Agency, website, http://www.epa.gov/risk/basicinformation.htm, accessed 8/2/09.

[32] I. Linkov, G. Kiker, R. Wenning, Environmental security in harbors and coastal areas: Management using comparative risk assessment and multi-criteria decision analysis, Springer, Amsterdam, 2007.

[33] J. Keisler, W. Buehring, P. McLaughlin, M. Robershotte, and R. Whitfield, Allocating contractor risks in the Hanford waste cleanup, Interfaces 34:3 (2004), 180-190.

[34] R. L. Dillon, M.E. Pate-Cornell, S. D. Guikema, Programmatic risk analysis for critical engineering systems under tight resource constraints, Operations Research 51:3 (2003), 354-370.

[35] M.E. Pate-Cornell, and R.L. Dillon. "The respective roles of risk and decision analyses in decision support. " Decision Analysis. 3:4 (2006), 1-13.

[36] M.E. Pate-Cornell, R.L. Dillon, and S.D. Guikema. "On the Limitations of Redundancies in the Improvement of System Reliability: The Case of the Mars Rovers." *Risk Analysis* (2004).

[37] A. P. Sage, Systems Engineering, Wiley, New York 1992

[38] G.S. Parnell, P.J. Driscoll, D.L. Henderson, Decision making in systems engineering and management, Wiley, New York, 2008.

[39] G. Richardson, and A. Pugh. Introduction to system dynamics modeling. Pegasus Communications, Waltham, MA, 1981.

[40] D. Duffie, and J. Pan, An overview of value at risk, Journal of Derivatives 4:3 (1997), 7-49.

[41] R. L. Dillon, and M.E. Pate-Cornell, APRAM: An advanced programmatic risk analysis method, International Journal of Technology, Policy and Management 1:1 (2001) 47 – 65.

[42] I. Linkov, F.K. Satterstrom, G. Kiker, C. Batchelor, T. Bridges, and E. Ferguson, From comparative risk assessment to multi-criteria decision analysis and adaptive management: Recent developments and applications. Environment International 32 (2006), 1072-1093.

[43] J. H. Lambert, K.A. Peterson, N.N. Joshi. Synthesis of quantitative and qualitative evidence for risk-based analysis of highway projects, Accident Analysis and Prevention 38 (2006), 925-935.

[44] J.H. Lambert, N.N. Joshi Diversification modeling for risk management of safety and security programs, Proceedings of the Ninth International Conference on Probabilistic Safety Assessment and Management (IAPSAM), Hong Kong, 2008.

[45] Department of the Army, U.S. Army Corps of Engineers, Screening for portfolio risk analysis (SPRA) guidelines for dams, Circular, No. 1110-2-xxx, December 2007.

[46] Federal Emergency Management Agency, FEMA: Risk Prioritization Tool for Dams Users Manual, FEMA P-713CD (CD-ROM), June 2008.

[47] A. Law, and W. D. Kelton, Simulation modeling and analysis, McGraw Hill, New York, 1999

[48] D. A. Mauney, Using financial/decision analysis risk optimization to prioritize maintenance expenditures, justify maintenance spending and maximise the NPV savings of the maintenance function,

Structural Integrity Associates technical paper, http://www.structint.com/tekbrefs/t99018r0.pdf, accessed 08/05/2009.

[49] C. Stummer, and K. Heidenberger, Interactive R&D portfolio analysis with project interdependencies and time profiles of multiple objectives, IEEE Transactions on Engineering Management 50 (2003) 175-183.

[50] K. J. Arrow, A difficulty in the concept of social welfare, Journal of Political Economy 58:4 (1950), 328-346.

[51] F. Hillier, and G. Lieberman, Introduction to operations research, McGraw-Hill, New York, 2005.