11. Modeling Business Value

Getting a Handle on the Intangibles of IT

In chapter 7 I argue that business value modeling and measurement must be a core capability of IT in order for IT to be able to articulate its concerns to the business. In this chapter I provide concepts and techniques for accomplishing this.

The failure to establish a common vocabulary and conceptual model for the business value of IT concerns, in a manner that has sincere participation and acceptance across both IT and the various business units, is a major root cause of the poor communication and disconnectedness that exists between business and IT in most large organizations. Yes, modeling the business value of IT concerns is hard, but it *must be achieved* if there is to be any hope of integrating IT planning with business planning. The implication is that IT leadership staff at all levels need to be versed in more than technology: they need to understand basic economic analysis. This book will not try to provide a complete treatment of the economic analysis of IT projects or issues, but I will describe some innovative ways of applying standard techniques, as well as some recently developed techniques, to model some of the hardest yet most important aspects of IT value.

Measuring the Value Of Architecture Choices

IT architects do not communicate well with managers in lines of business. They might have a good relationship, but if an architect is asked about the concerns that they deal with, they will mention things such as "technical risk", "coupling", "encapsulation", "reusability", the need to "refactor", and "transactional integrity". These terms mean little to people in business; yet these terms mean a great deal to IT architects and they are extremely important concepts that have a significant impact on business agility and lifecycle cost. Why then is the communication between architects and business people so difficult?

IT architects are not trained to relate architecture concepts to tangible business value.

The problem is that IT architects are not trained to relate their concepts to tangible business value. For example, "coupling" refers to the amount of interdependence between components. This impacts business agility because a high degree of interdependence (coupling) generally means that changing either component is costly and difficult. Yet, architects do not have standard practices for expressing coupling in terms of lifecycle cost or agility.

This situation is the fault of our educational system. IT architects need to provide a link between technology and business, and as such they should be versed not only in how to express decisions in technical terms but how to express them in business terms as well. Architects need to have a repertoire of templates for modeling the business value of the technical concerns that they know are important.

Despite the fact that this is seldom practiced, it is very possible to do this. Later we will see how to model the value of agility by comparing possible future scenarios. (See page 259.) Coupling impacts agility and so it can be modeled in the same way. Technical risk can be modeled as negative opportunity.[1] (See Modeling Risk on page 235.) Refactoring refers to the simplification and internal reorganization of a component in order to enhance its reusability and maintainability, and so the value of refactoring derives from increased agility, a higher level of reuse, and lower lifecycle cost. Other architectural principles can be modeled using similar techniques. Thus, we see that the excuse that architects do not normally model business value does not hold water. They could, and they should – if they knew how. Since they are not trained to do so, you will need to team them with staff or consultants who

[1] In this book, and in most literature in the security community, a threat is something bad that can happen, and risk represents the potential loss due to that threat. There are other definitions of risk and threat though. E.g., the Project Management Institute PMBOK methodology inverts these definitions.

have experience modeling business *value* and challenge them to create models of their architecture principles.

Business value is very different from cost.

It is important to realize that business value is very different from cost. There is a great deal of literature that provides techniques for cost estimation in IT projects. Cost only impacts value when the value derives from reduced cost. Other kinds of value such as increased agility cannot be measured using cost-based techniques.

Value should be expressed in terms of goals. At the highest level, this is usually shareholder value, but can be other things. At intermediate and lower levels it can be objectives having to do with operational efficiency, revenue, data quality, or other success criteria. Thus, increased agility can translate into the ability to introduce more products as the need arises; this value need not be expressed in dollars, but can be expressed in terms of the months to introduce a new product. If that has meaning to the business and rolls up to shareholder value in the corporate plan, then there is no need to express the value in dollar terms in a departmental plan.

The Components of Business Value

When measuring business value one must establish a planning horizon. That is, value needs to be maximized over a particular timeframe. As discussed earlier in chapter 7, the time frame makes a large difference. For example, actions that accept long-term risks might be optimal if the time frame in question is one year, but other actions that manage those risks better might be optimal if the planning time frame is five years.

Organizations often have multiple planning horizons, such as short-term budget cycle planning and longer-term roadmap planning. Longer-term objectives and opportunities should be offset by longer-term risks and lifecycle costs. Shorter-term objectives and opportunities should be rolled up into longer-term objectives to make sure that shorter-term tactics do not compromise longer-term objectives.

The overall business planning horizon should be long enough to span the lifecycle of major systems. In Making Technology Investments Profitable, Jack Keen and Bonnie Digrius say pointedly, "Benefits come true when management monitors every step of the way, from project funding through system retirement."[1]

This is quite a bit longer than the span over which most organizations plan – in fast-paced organizations roadmaps are often no more than three years out at best. However, one does not need to actually *plan* over that period: one merely needs to account for costs and business impact effects that manifest over that period.

The expected value of a business capability over a planning horizon is composed of:

1. The expected income from the foreseen opportunities that the new capability will enable, *minus*
2. The expected investment, *minus*
3. The expected lifecycle cost, *plus*
4. The expected value of *unforeseen* future opportunities that the new system will enable the organization to capture.

These terms should all be adjusted to account for the passage of time, using a "net present value" (NPV) formula, but I am not even going to account for that, because as it turns out, adjustment for time tends to be a small affect compared to other factors in IT decisions, and I do not want to distract from the core concepts here by presenting complicated formulas.[2]

The first two items in the list above are straightforward: the opportunity and the investment. There are the two main aspects to any standard business plan. However, the standard type of analysis has a deep flaw: it is deterministic. That is, it assumes that the outcome is certain. A deterministic approach is fine for many investment situations in which risk is treated separately for each investment. However, in a large organization in which investments in many activities occur continuously and all impact each other on an ongoing basis over time, the business model must incorporate the inherently non-deterministic nature of investment into the core value model.

Consider Figure 13. (Note that the vertical and horizontal value scales are different.) When a new capability is contemplated it does not yet exist. Since there is a chance (often a large chance) that the

---

[1] Ref. [Keen03], page 159.

[2] For those who wish to do a more precise calculation, it is almost trivial to define the sources of value as a time series, dividing each term by a discount rate, using a standard NPV approach. However, the discount rate used should represent the value of money, and not attempt to account for risk.

capability will not be completed and deployed, and even if it is deployed, might not live up to expectations. Thus, at the outset of planning the expected value of the future capability is substantially less than its potential. It is not zero however, since if it were there would be no point in even starting the effort.

As work on a capability proceeds, the chance of ultimate success normally increases. It may actually decrease if troubles arise that indicate that the capability might not be feasible. In some cases, new information (such as development failures) might cause the chance of success to drop below the point at which investment is worthwhile. Regardless, the expected market value is the product of the probability of ultimate success and the estimated maximum market potential. This is shown by the curve in the figure.

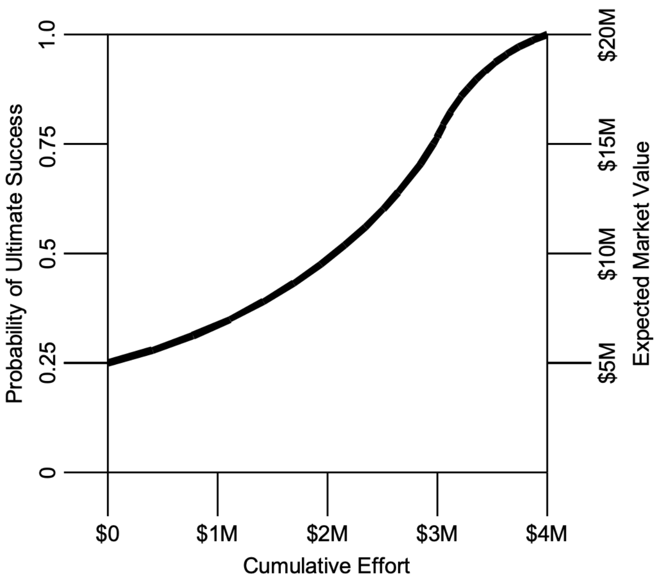

Figure 13: Expected market value approaches its maximum as work proceeds and the probability of successful completion reaches 1.

It is noteworthy that the curve in the figure takes on kind of an S-shape, characterized by a gradual – almost horizontal path – followed by a steep climb, and then a gradual approach to a probability of 1.0. The initial gradual ascent represents the period during which "show stopper" concerns are addressed. These are often technical challenges which are best dealt with by early proof-of-concept efforts.

In an IT project one usually does not consider what happens in the marketplace after deployment: that is considered to be the problem of the business side. However, to model the value of a capability realistically, one should account for all risks, including market failure. If that is not possible, then IT has no choice but to assume the market potential to be that which is articulated by the business.

Now consider the last item of the four above, the value of unforeseen future opportunities. This item represents the business value of business agility. Agility is the ability to respond to new opportunities as they occur. The specific opportunities are unforeseen – otherwise one could plan for them and one would not need to respond to them as they occur. Thus, having agility means having an ability to react in ways that one cannot predict. To preserve agility one must invest in preserving agility, because it takes effort and forethought to preserve agility. Apart from the cost though, the expected *value* of agility is not related to its cost: its value is the value of the opportunities that can normally be expected to come along.

The third item above, the expected lifecycle cost, is often treated as the cost of the expenditures that go into building and maintaining a system. The true costs go way beyond that though. If one is to use lifecycle cost to estimate the impact of changes to a system, then one must estimate a true lifecycle cost that accounts for all of the indirect impacts, including impact on the *cost of reacting to future opportunities*. In other words, we must account for the cost of foregone agility. We must also account for the expected cost of failures that occur due to risks coming true as a result of inadequate investment in risk management. These components are the cost part of the equation, in contrast to the opportunity value discussed in the previous paragraph. Any business plan must account for opportunity value and investment cost. In the sections follow I will break these components down further.

Accounting For the Value of Agility: Making the Case for Agility

Chapter 7 explained how business agility relates to many IT architecture principles, and makes the point that when one applies these principles, one must articulate their value in business terms. Here I want to explore an approach for doing that. Later in this chapter I will revisit this topic and actually work through an example.

If you expect your IT staff to preserve business agility, then you must be willing to measure the long-term impact of a series of actions that individually do not significantly impact agility but in the aggregate destroy it. This means that initiatives must anticipate the impact of any deviations from the intended architecture, and that problems that occur later must be traceable to the decision-makers who caused those problems by allowing rules to be relaxed in the interest of expediency. This is a tall order, but it will occur naturally once the organization starts to plan based on business value models, as those plans and measurements will accrue over time.

If you expect to preserve business agility, then you must be willing to measure and trace the long-term impact of a series of actions.

Business-focused architects (which all architects should be) should be able to produce a *quantitative* estimate of the impact on expected lifecycle cost (item 3 above) of any major design decision. Thus, when there is a call to compromise the architecture of a system, an architect should be able to estimate what the business impact will be in terms of increased lifecycle cost.

Item 4 in the list of business value components is the value of agility: the value of being able to react to change and capture new, *unexpected opportunities that were not planned for*. Measuring the value of agility is much harder to conquer. However, it is perhaps the most important value of all. It is impossible to know the value of opportunities that have not yet shown themselves, yet this is what agility is for: to be positioned to seize such opportunities. Therefore, they must have value, and you must put some kind of metric on them if you are to be able to trade off agility with short-term cost savings.

While one cannot predict actual future events, it is often possible to anticipate *classes* of future events.

Note that while one cannot predict actual future events, it is often possible to anticipate *classes* of future events and estimate their relative probability, based on prior history. For example, while one cannot predict that a security failure will occur – with all of its associated costs in lost customer confidence – one can nevertheless anticipate that certain classes of failure will occur, and project their likelihood over time given past history combined with industry statistics and knowledge of the organization's level of precautions. Similarly, while one cannot predict that the organization will decide to introduce a particular kind of new product within the next two years, one can look back at the history of the organization to see how often it has historically introduced new products, and use this as a benchmark for what is likely to occur in the future.

The method known as *Real Options Analysis* (ROA) is an analytical approach to addressing uncertainty and can be used to measure the value of agility. The real options method is explained more later, but a simpler approach that is usually adequate for our purposes is presented in the sections that follow.

It is not always necessary to estimate the quantitative value of agility: one often only needs to qualitatively estimate the relative agility afforded by various choices – and recognize when agility will be whittled away as a result of a particular choice. However, when making a business case that *trades off agility for other concerns*, such as cost, immediate opportunities, and risk, *one must assign a value to future agility*: otherwise, there is no basis for comparing various strategies that each produce a different mix of these attributes. The value metric, and the methodology for producing it, must have consensus across IT and the organization's other operating units. Otherwise it will not be trusted.

History can be a guide to the value of agility. For example, if your organization is a product company, and one looks at product plans from five years earlier, one can ask introspective questions such as, *Were there any times when we were not able to react as quickly as we would have liked, and what would have been the value of being able to react quickly?* Asking these kinds of questions can be used to establish a benchmark for the prior value of agility, and be a basis for reaching consensus on an educated guess for its future value.

Consensus is important, because estimating the value of agility is at best educated guesswork, but if such a measure is to be part of a business model that is used to by IT in proposals that compare the relative values of architecture features in IT plans, then the measure needs to have the credibility that only a consensus-based process can produce.

Modeling Risks: Making the Case for Risk Mitigation

In chapter 7 I present a way of measuring the business value of system features based on their contribution to reducing the risk of realizing a future business capability. In this section I want to discuss risk in general, including exogenous risks that require mitigation to some degree.

Risk is often addressed in initiative proposals by simply describing the risks. While this type of intangible approach at least provides for the identification of risks so that attention can be paid to them, it does not enable tradeoff analysis with respect to opportunity. After all, risks are sometimes worthwhile if the opportunity is great. Further, different organizations have different levels of risk tolerance, depending on the expected longevity of the organization and planning horizon of the organization – i.e., how impatient its investors are.

For many kinds of risk a mere identification of risk is appropriate, but for IT the risk of reliability-related failure and the risk of security-related failure need to be compared; otherwise one might spend too much on one and too little on another. Further, these need to be compared with other sources of risk so that tradeoffs can be made intelligently. More importantly, the tradeoffs need to be explicit so that the tradeoff criteria can be reapplied and adjusted if necessary as the situation evolves and as new proposals crop up.

The real question with risk management is how much to spend on mitigation, and is the expense such that the opportunity is not a good investment? One therefore needs to create a framework for determining the optimal amount to spend on mitigation for various levels of opportunity. Such a framework makes investment in risk management transparent just as other kinds of investment are.

A risk mitigation model also enables one to provide a credible estimate for the level of risk that the organization has and the investment required to attain it. This type of information would be strong supporting evidence in any type of audit situation. Rather than telling an auditor, "we are doing all we can to reduce our risk", one can say, "our risk will be X based on our investment in mitigation, according to our model that links that investment to the level of risk".

Creating such a framework is relatively straightforward. I have done it for reliability risk mitigation and for security risk mitigation. A detailed treatment is beyond the scope of this book, but I will provide a sketch. The basic approach is to model risk as a detractor from the *expected return* on an opportunity. The impact on the risk for various expenditures on mitigation must be plotted and a maximum found for the opportunity's expected value. That is, the expected value must subtract the cost of mitigation as well as the expected losses due to some risks coming true and optimized over the various possible levels of mitigation. Once the optimum is found it can be adjusted manually based on how risk-averse the organization is.

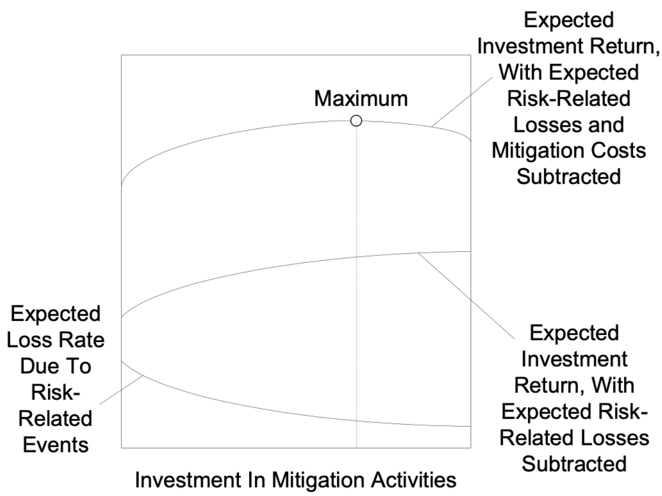

Figure 14: Plotting expected return of an investment opportunity with risk and mitigation included.

Figure 14 illustrates this technique. This figure provides a risk tradeoff analysis for a particular investment opportunity. In this figure the lowest curve plots the expected loss incurred as a result of risks coming true according to their expected rate of occurrence for various levels of expenditure toward the mitigation of those risks (including the zero level, which represents the inherent or baseline risk). The middle curve plots the expected return of the opportunity with the expected losses subtracted. The uppermost curve (offset to appear above the others for readability) also plots the expected return but subtracts both the expected losses and the cost of mitigation, providing a total picture of the investment. If the total amount of investment available is a fixed budgeted amount, then the uppermost curve must incorporate the fact that investment in mitigation subtracts from investment in the opportunity itself – i.e., safety costs money. In any case, the point of maximum on this curve indicates the point at which *the expected value of the total investment is maximized* as a function of spending on risk mitigation. If an organization is risk-averse, then the risk represented by the point of maximum might still be too high, and the organization may choose to reduce risk to a predetermined level or apply a weight that is an inverse function of risk.

The process of calculating these curves is pretty straightforward economic analysis and the basic process is illustrated in Figure 15.

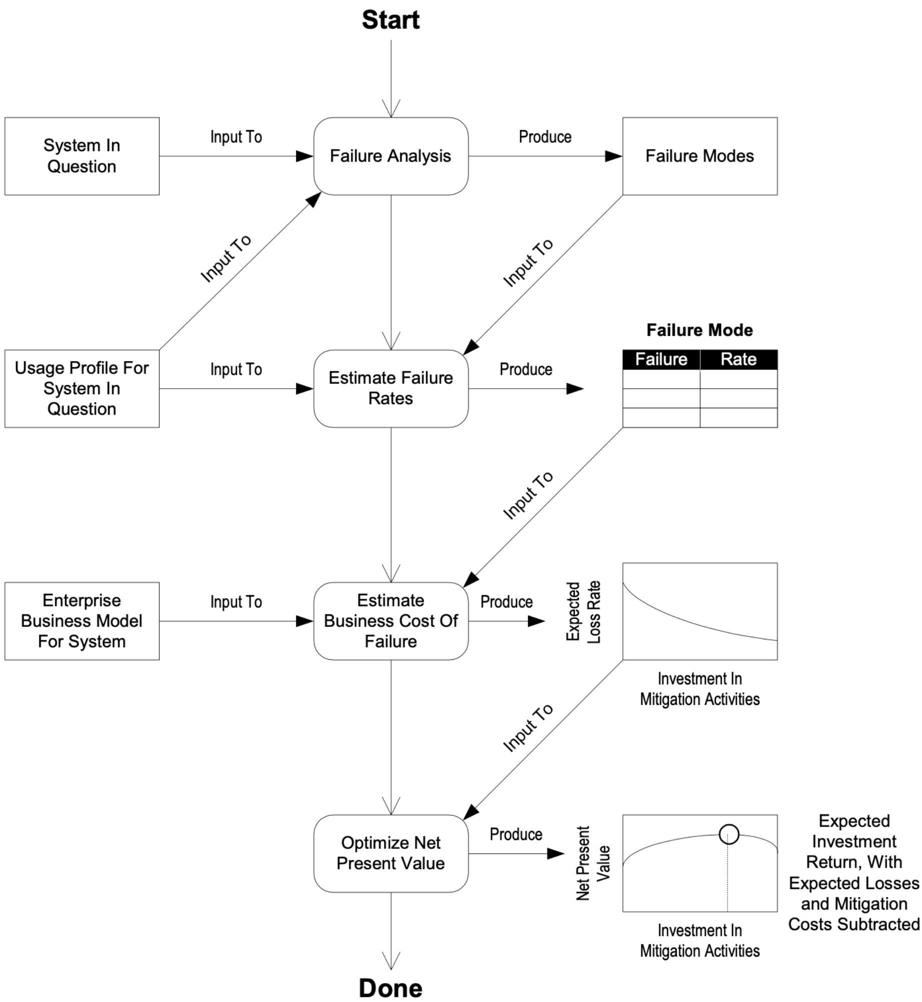

Figure 15: Combining failure mode analysis with value optimization to determine the required level of risk mitigation.

The challenge here is that risk must be quantified as an expected shareholder value[1] loss rate over time. In the figure, this is represented by the "Estimate Business Cost Of Failure" step. In doing so one must consider the various failure modes – the ways in which failure can occur – and their *tangible* impact on the business, including loss of brand value, loss of future business, and even losses due to litigation when risks occasionally come true.

Evaluate Strategies, Not Isolated Decisions

Traditional gap analysis presumes the existence of a definition of a future state, and evaluating a decision amounts to computing the ROI of that decision compared to a competing decision.

The dilemma is that, as has been pointed out earlier, one never arrives at the future state: the future is almost always different than the way you predict it will be. Therefore, what credibility does gap analysis have? And further, what legitimacy does the ROI analysis have if the future is so indeterminate?

A cornerstone concept of newer financial methodologies such as Real Options Analysis is that it makes no sense to compute ROI based on a particular future state: rather, one should assume a probabilistic distribution of future states. For example, rather than assuming that the actual market demand for a product will be some number X, project the expected distribution of the demand, with a mean and a standard distribution on the probability density. Thus, one might say that the expected demand will be X, with a standard deviation of 0.3 X.

With such uncertainty, it also does not make sense to predict that one will have certain choices to make in the future, and so a single decision today will be followed by choices that are unknown, all of which will impact the return resulting from the choice made today. In other words, projecting a return on today's decision is somewhat

[1] There are many approaches to estimating shareholder value. The intention here is to use an approach that is aligned with the intentions of the shareholders: if the intention is to liquidate the firm after a short-term rise in value (oil drilling ventures are often done this way), then value should be based on the liquidation value. If the intention is to hold shares for a protracted period and the firm is a public company, then value should be based on factors that affect the traded price of the company such as residual income from investments, reputation, debt ratio, and so on.

foolhearty, because events might occur that cause the organization to change direction and revise today's decision.

Instead, it makes sense to evaluate the return from a particular *strategy*, rather than a particular decision. Consider a situation in which IT wishes to compute the return from building flexibility into a new business system. IT might compute the return from that decision to be R, based on the assumption that that flexibility will actually be utilized in the future. However, it is likely that the future will bring more decisions, and those decisions might involve tradeoffs of flexibility versus expediency. What then is the return, over some period, of consistently choosing the more flexible choice, each time a decision of that kind presents itself?

That is, if one always chooses the more flexible path, will a series of decisions of that kind pay off eventually? That strategy of choosing flexibility has certain ROI characteristics, and to evaluate them one must model the generic choice of choosing flexibility, and evaluate what would happen as the scenario plays out again and again over time.

This approach is much more effective for modeling the value of many kinds of IT strategies, because many IT strategies have to do with preserving flexibility or maintainability, and the value of such a choice only accrues over time and over the course of many decisions. This approach is illustrated later in this chapter.

One aspect of modeling the value of a strategy is that the expected value of the business should accrue continuously. That is because if one pursues a strategy that has long-term tangible value, then at any point the expectation of realizing the long term value increases, and so the business systems for realizing that value are valued higher. It is as if you have planted a seed, and grown a fruit tree half way to the point of bearing fruit, and so the tree has market value ("economic value") because it *will* bear fruit, even though it is not yet bearing fruit, and as you nurture that tree and continue to invest in it, its expected value increases continuously.

Be Careful Not to Double-Count Benefits

Management, and portfolio boards in particular, often want to know the expected business benefit from each element of a plan, or each portfolio line item. Unfortunately, items are often inter-linked.

This is another reason why it is important to express the value of strategies, and not particular decisions. Interdependence of value is a good thing: it is a sign that efforts are inter-linked and working together. It means though, that a single element of a plan cannot claim the value of the plan. For example, assume that a new strategy involves creating a data management group, and to create that group, it is assumed that a data governance process will be established, and that the value of the better data quality that will result is $10M per year. When asked the value of the data management group, there is a tendency to say "$10M per year", and when asked the value of the data governance process, there is also a tendency to say "$10M per year". Management will then be confused and add these numbers and get $20M per year.

To be clear, rather than creating a large table of dependencies, it is better to express the *strategy* of data management, which requires subordinate investments in data governance, among other things, and give this strategy a value. When asked about individual elements of the plan, they do not have independent value, so do not give them one. If management says that you can only have some of the plan's elements but not all, due to funding constraints, then you have failed to give management the right choices. If the plan's elements are interdependent, the it is your job to present options that include a scaled down version of the plan, perhaps one that unfolds more slowly, but that still has the elements that it needs to succeed.

An Example

Let's look at a simple example to clarify all of this. Again, this book is not intended to be a text on economic analysis of IT projects, but the techniques presented here are only now finding their way into mainstream practices for IT, and so a short example is in order.

In this example the cost of money over time is omitted. This is a huge simplification because the cost of money is extremely relevant in a real analysis, but it is not relevant to the novel techniques that I am trying to explain.

Consider a situation in which an opportunity presents itself to implement and market a new service. The market opportunity is estimated to be at least $12, but the degree to which that can be realized depends on the level of investment, as shown in Table 12. The table does not include the downside: risks and the potential costs that they represent.

Table 12: Market opportunity realization as a function of investment in marketing.

Investment In Marketing	Expected Opportunity Value (Revenue – Direct Costs), Assuming No Failures
$1M	$6M
$2M	$11M
$3M	$14M
$4M	$15M
$10	$18M

From a risk perspective, the possible failures that are anticipated include: (a) some sales are returned due to defects; (b) financial injuries attributable to the service. These risks are summarized in Table 13.

Table 13: Failure modes, i.e., risks, and their expected long-term average cost.

Failure Type Name	Nature of Losses	Expected Cost To Enterprise Per Incident
A	Restitution to affected consumer customers.	$0.1M
B	Loss of reputation and future business. Potential loss of some commercial customers. Legal costs.	$50M

To reduce these risks, the following mitigation strategies considered: (1) a range of quality control measures; and (2) a range of intervention measures designed to avert financial injury to the consumer in most cases. If these strategies are implemented, the consequent expected rate of occurrence of each type of risk is summarized in Table 14. The mitigation strategies are assumed to be independent. If they were not, then Table 14 would need to list a set of technically viable scenarios that represent mixtures of investment in each mitigation strategy.

The total capital budget for the project is $10M, and so the investment in marketing and mitigation must be balanced and hopefully optimized.

Table 14: Expected failure rates for each risk type, for a given investment in mitigation strategies.

Failure Type	Mitigation Investment	Expected Failure Rate Over Planning Horizon	Incident Cost
A	$0	50 incidents	$.1Mx50=$5M
A	$1M	5 incidents	$0.5M

Failure Type	Mitigation Investment	Expected Failure Rate Over Planning Horizon	Incident Cost
A	$2M	3 incidents	$0.3M
A	$10M	~0.01 incidents	$1K
B	$0	2 incidents	$100M
B	$1.5M	0.1 incidents	$5M
B	$2.1M	0.01 incidents	$0.5M
B	$10M	~.0001 incidents	$5K

Analysis

If we plot on a graph the expected revenue versus investment in marketing, and the expected losses as a function of investment in mitigation, we get the graphs shown in Figure 16.

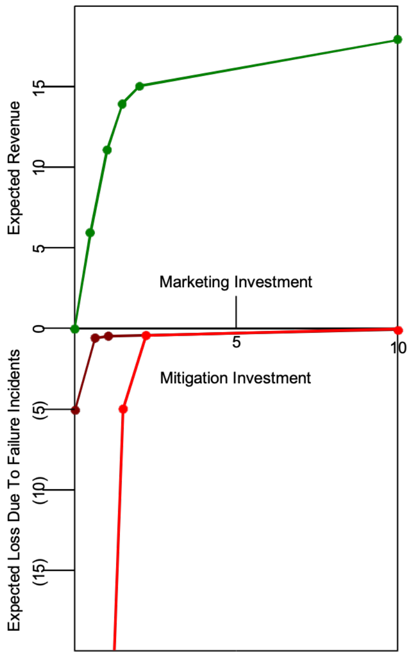

Figure 16: Revenue and losses as functions of investment in marketing and risk mitigation, respectively.

Clearly risk mitigation is important in this case: in the extreme scenario of no investment in mitigation, the potential loss to the company, due to lost future business across all product and service lines, runs way down off the graph, overwhelming any revenue from the sales of the new service. On the other hand, it would not make sense to invest the entire project budget in mitigation, since then there would be no point in investing! Clearly, investment in marketing and mitigation represent a zero-sum game, with an optimal mix of investment in marketing and mitigation somewhere in the middle.

It is worth noting that in most (but not all) organizations the decision on the mix of IT investment choices is made in a fairly ad-hoc manner: technical staff might be asked to present a couple of mitigation scenarios, or in many cases they are expected to figure this out on their own and submit a budget for development with appropriate mitigation implied. The problem with this approach is that in many cases the tradeoffs are not straightforward, and the business needs to be involved in deciding that tradeoff. To do this they need to know the business value of the various mitigation strategies, expressed in terms that they can understand, either in terms of dollars, or in terms of a risk scale that they understand and can related to priorities and ultimately dollars or enterprise value.

Finding the optimum mix of investments is of course an optimization problem and the most practical approach is often to simply define different scenarios and compare them. We are trying to maximize the present value of:

> Revenue – Direct Costs – Incident Costs – Mitigation A Costs – Mitigation B Costs

where each term is an expected value, since we can only estimate each. Comparing scenarios is actually more practical in most cases, because usually mitigation strategies are interdependent. This is especially true in the security realm, where a particular strategy might serve to mitigate more than one type of risk. Thus, we end up with a table of investment mixes and their expected costs and revenue, and plotting the formula above results in a graph something like Figure 14.

Analyses Of Risk And Opportunity Should Be Integrated

In chapter 4 I refer to a poll taken by the Northern Virginia Java User's Group (NovaJUG) to assess how widespread risk management practices are in Java software projects (see page 42). One of the more astute members of the NovaJUG responded to the survey with an anecdote about how he has practiced risk management in prior projects. According to him,

"…we would come up with a risk, write it down, then rank on a scale of 1-10 the impact it would have on the project and a likelihood it would happen. We would multiply these together, rank them all, and make sure we had some kind of mitigation strategy (to prevent it from happening), or plan B, if it did happen."

In other words, he used an approach that modeled risk according to the product of the probability of failure and the cost of failure, and then used this result to rank (prioritize) efforts to devise and implement mitigation strategies.

This is indeed a viable approach to modeling risk for the purpose of ranking. One sometimes only needs to rank risk to determine where next to invest effort in mitigation. However, this approach is inadequate if one must trade off the business value of mitigation against the business value of adding additional features.

For this reason it is important that risk and business value models are integrated. The drivers for accounting-related risk mitigation are often auditor-related but they should nevertheless not be dealt with independently. That said, organizations often need to have a focused risk mitigation effort in order to get through an auditor-related crisis, since auditors might not appreciate that accounting accuracy is not the company's only concern. Therefore, achieving integration between risk and value management are a great challenge. On the other hand, auditors are primarily concerns with accuracy, and accuracy is the friend of both risk and value management.

This means that risk management needs to work with the business and explain the framework to auditors. Transparency is key. Further, risk and value both need to be incorporated into architectural analysis in order to make the architectural analysis "real". If this is done, it will be clear that accuracy is a way to achieve business goals related to both risk and value, since accuracy leads to better measurement and therefore better business choices.

Triage First

When planning a risk management strategy, the best approach is to start with a two-pronged approach that includes "triage" of existing known sources of risk, and develop in parallel a business value-driven maturity framework that incorporates both risk and business value into planning. Maturity considerations are discussed further on page 265.

I have used the triage approach many times to relieve a crisis. A crisis presents an opportunity to explain root causes when attention is on an issue. For example, in the case of the Guru-Run Technology Provider, it was a crisis that precipitated my involvement in addressing reliability problems with the system, and while the root causes were clearly systemic, it was critical to maintain the confidence that the customer had in the contractor and get the system into a state in which it was recovering automatically from failure and not losing transactions.

One should leverage the attention that a crisis brings to establish expectations and consensus about the business value model for a long-term strategy that addresses the root causes. Once the immediate concerns are alleviated through the triage interest will fade, but the business value model can be used for future proposals and it will be understood. Risk triage solutions are not durable and do not enable agility, maintainability, or transparency, and so the business value model for elevating maturity must focus on those aspects. We did this for the Guru-Run Technology Provider, and I recently followed up with this project and learned that the system ended up being re-written from scratch using an entirely different set of technologies and with much more attention to reliability from the outset.

Features That Have No Independent Business Value

New features can always be developed for automated systems. However, as explained in the section Forget The "End State": Design A Capability For Evolving The Current State In The Right Direction (page 72), automated systems are not by themselves business capabilities, and a *business capability* cannot be claimed to have been created until it is demonstrable. This is because the IT components of a capability are merely the automation aspects of a capability and by themselves are useless. One could therefore make an argument that the only IT features that have independent tangible business value all by themselves are those that can be immediately deployed with no user training.

In the section Measure Future Value At Risk (page 123) I explained how one can decouple the values of interrelated features of a system and relate them to the overall value of a business capability. The justification for adding a feature that has no independent business value is to invest in the opportunity to build toward a future capability. I will refer to such features with no independent value as *precursor features* because they are precursors to the completion of a complete capability that has tangible business value. (In Figure 8 on page 126, feature set A represents a precursor features with respect to the business capability.) Having "tangible business value" means that the business value is represented in some manner in a model that has the consensus of executives.

A business capability cannot be claimed to have been implemented until it is operational.

Precursor features represent partial progress toward a business capability. The progress is real, but it has no usefulness to the business. Progress on such features are analogous to the progress that one makes when one tears down barriers that stand in the way of a goal. The barriers are not themselves part of the goal, but removing them is nevertheless progress.

Unlike a precursor feature, a business capability cannot be claimed to have been developed until it is operational. That is, the business processes – human and automated – that implement the capability must be up to speed and running. The hardest part of creating a new capability is the last "five percent" – which can turn out to be 50% – and a good part of that last percent is the work involved in elevating the skills and understanding of the staff and building a shared understanding of how things will work at all levels. Setting the bar of what constitutes a true capability therefore must include these cross-the-finish-line efforts to ensure that metrics for progress are not misleading.

Estimating the Value Of Precursor Features

Precursor features have no independent business value, but demonstration of progress can be in terms of a technical test suite that includes certification for compliance with all enterprise concerns. For features that have actual end-user functionality, the demonstration should be both in terms of a test suite and also in terms of end-user acceptance.

Precursor tasks that have no independent business value nevertheless increase the *expected value* of the end state.

Any future capability is tentative and represents risk in terms of the following factors:

1. The successful completion of the capability; and
2. The validity and success of the value proposition of the capability, which is merely theoretical until the capability is put into use. I.e., do forecasts of upcoming opportunities pan out?[1]

Precursor features or tasks are an investment in a planned future capability. Once completed, such features or tasks represent increased progress toward the planned capability but do not represent capabilities themselves. As such, they represent value because this progress has value. The value of this progress derives from the following elements:

1. Reduced risk in realizing the opportunity, and
2. Reduced remaining investment is less for completing a capability.

Risk is reduced because the feature represents successfully completed work toward the intended capability. The completion of this work eliminates the risk that the work might not be completed or that it might not work – that is why capabilities cannot be considered to be complete until they are demonstrable. The second element, reduced remaining investment, results from the fact that the work required to complete the capability has been reduced as a result of the completion of the feature. Thus, the value of the feature does not derive from the effort put into it: rather, it derives from the fact that the business case calculated after the feature has been completed is better than the business case calculated before the feature was completed.

The reduced risk (element 1 above) should directly increase the *expected value* of the planned future capability that the feature is an element of. Since any future capability is tentative and represents risk in terms of the successful completion of the capability, any progress toward completing the capability represents reduced risk

[1] This is sometimes referred to as the "option value of investment".

and therefore increases the expected value of the capability. This is a purely probabilistic effect.[1]

The reduced remaining investment (element 2 above) means that the investment required to complete the capability is now less since the feature has been completed and tested. This progress results in a new, improved business case from this point onward: one should calculate the work required to complete the capability, and compare that work with the work that had been estimated before the feature was developed. The change represents the reduction in planned investment, given that the feature has been successfully developed.

You might wonder, what about the actual effort that was put into developing the feature?[2] Unfortunately, that is not relevant to calculating the *value* of the feature. The feature might have been implemented in a very efficient and timely way, and that efficiency has value because it reduces the cost of capability development. However, the cost of the feature needs to be accounted for in terms of the actual investment in the capability. *It is not part of the feature's value*. Rather, it is what is known as a "sunk" cost – money spent and gone forever.

Let's briefly consider a simple example. Suppose that a capability will require several features, and before a feature is developed, the project manager estimates that the work required to complete the full *capability* is ten person-weeks. Work on the first feature is begun, and the feature is completed in one person-week. The project manager then estimates that the work required to complete the capability with the new feature incorporated will be seven person-weeks. Thus, the project has moved ahead by three person-weeks as a result of the new feature; yet, only one person-week was invested. The fact that one person week was invested has no bearing on how much time or effort remains in the project: all that matters is what will be required to complete the project, irrespective of how much effort has already been invested – wisely or otherwise. The work put into the feature is an investment, not a value.

To summarize, the value received after adding a precursor feature is therefore:

---

[1] For more discussion of this type of analysis, see [McAfeePr06], §4.3.3, Investment Under Uncertainty.

[2] The amount already invested is often referred to as the "sunk cost" – that is, the amount that has already been sunk into the project.

1. The increase in the expected value of the opportunity (due to reduced risk of completion), *plus*
2. The reduction in the remaining investment as a result of the existence of the added feature.

The threats that exist include (1) that the value proposition is miscalculated; and that (2) the project might not complete as planned.

The expected value E(V) of a capability is the expected value of the opportunity (O), adjusted for the risk (R) of failing to realize the opportunity, minus the investment cost. (One can also add to R other risks or possible costs that are of concern.) The risk of failing to realize the opportunity is primarily impacted by the risk of failing to complete the project successfully. This, it essentially represents the development risk, in terms of the potential lost opportunity.

In mathematical terms,[1]

$$E(V) = E(O) - E(R) - E(C)$$

Let's then consider a situation in which we are building a capability, a feature at a time, and we wish to calculate the value of a particular feature. At iteration 1 the expected value of the system (incomplete capability) is:

$$E(V1) = E(O) - E(R1) - E(C1)$$

and at iteration, after another feature has been completed, two its value is

$$E(V2) = E(O) - E(R2) - E(C2)$$

The change in value as a result of the addition of the single feature is the difference between these:[2]

$$\begin{aligned} \Delta\, E(V) &= E(O) - E(R2) - E(C2) - [E(O) - E(R1) - E(C1)] \\ &= -\,\Delta\, E(R) - \Delta\, E(C) \end{aligned}$$

where

Δ E(R) is the change in the expected risk in completing the capability. The change period is over the time that it takes to develop the feature. Since this change should be negative

[1] Cost of money considerations – i.e., present value adjustments – are not included here for simplicity.

[2] We are assuming that the estimate of the business opportunity value does not change.

(corresponding to a decrease in risk), $-\Delta E(R)$ should be positive.

Δ E(C) is the change in the expected cost to complete and deploy the capability once the feature has been completed and integrated. Since this change should be negative (corresponding to a decrease in risk), $-\Delta E(C)$ should also be positive.

Both terms on the right should be positive, resulting in a net positive value for the new feature.

This situation is illustrated in Figure 17, which is based on Figure 13. In Figure 17 I have added the estimates given by IT for completion of a project at the start of the project and at the end of each of three periods of work. For example, at the start of the project IT estimated a cost of completion of \$3 million (considerably less than the actual final cost of \$4 million). At the end of period 3 the cost of completion from that point was estimated to be 0.9 million. (The cost-to-complete estimates made at various points during the project are shown as a thin gray line, and these points apply only to the monetary scale on the right side of the graph.)

Since in this example the actual cost did not track the projected cost, the value proposition for the project changed over time. This is shown by the thick gray line, which plots the expected business value (the expected market value minus the cost of completion). Progress from one work period to the next is simply the change in the height of the thick gray curve.

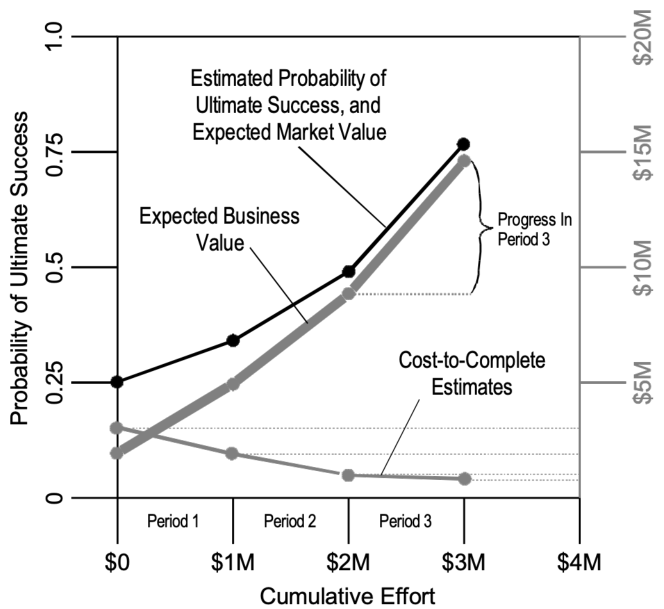

Figure 17: Expected business value changes throughout a project, and progress should be measured by the change in expected business value.

This approach provides a powerful instrument for projects that utilize agile development methods. In such projects, customers are allowed – encouraged – to reconsider requirements throughout a project. In that way the risk of building the "wrong system" is reduced. However, that risk is somewhat offset by the risk of constantly changing the system and getting nowhere – at great cost. A value-based analysis will show you clearly when you are spinning your wheels, because when earlier work is changed or discarded, your progress is slowed. Thus, instead of focusing only on cost and budget, one also looks at how close you are to creating real value.

It is important to note that progress is completely independent of the amount of money spent on work already done. Money spent this way is a sunk cost: it is irretrievable. Money spent on capital assets can often be retrieved, and that must be accounted for, but I will ignore that to keep things simple here.

It is also important to note that progress is not equivalent to a running ROI. An ROI can only be calculated once a capability is in operation and generating returns. An ROI calculation is a retrospective calculation, that is a function of the total amount of money spent. One could calculate a running projected ROI as a system is built, but that would be a different calculation. (See if you can derive the formula.)

Of course, in a real analysis, the time value of money should be accounted for by using time-discounted present values instead of fixed values, but I have ignored that to keep things simple and make my points. It is an important consideration however, because some kind of internal rate of return should be used to assess whether a project's value continues to justify that it continue: as the expected value fluctuates, it should be compared with the expected total investment over the time period. These detailed refinements are beyond the scope of this book, and indeed they are usually not significant for the time scales and large uncertainties that IT deals with, but I think you get the point.

There is one way in which it impacts work though: it is important to eliminate show-stopper obstacles early in a project, because doing so increases the slope of the risk curve early. In this way, project managers have more information early about whether to continue a project. In the aggregate over many projects, they must be able to decide which to continue and which to not, and getting onto the steep part of the risk curve early provides them with this information before making a large investment. Thus, from a *portfolio management perspective*, one should endeavor move forward (sooner) in time any tasks that will increase the probability of eventual success. This is common sense to a large degree, and it is consistent with the analysis here.

The question arises as to how one should estimate the probability of success. It might seem that this is a highly subjective measure, but it is important that your own subjectivity be removed from this calculation. The challenge – as with projecting costs – is to be as honest as possible. This is best achieved by compiling statistics on similar efforts within the organization and a set of criteria for predicting future success likelihood – i.e., a model. You can also ask "experts" to make educated guesses, as long as you calibrate those in some way against actual data (perhaps anecdotal) or an independent source. This is not rocket science: we are talking about a spreadsheet of estimates.

Another way of looking at project progress is based on the original estimated cost of features. Going back to our formulas, if we isolate our analysis to only the work done on a feature and ignore work done on other features, then our estimate for completion should be our original estimate with the estimated cost of the feature subtracted out:

$$E(C2) = E(C1) - C_{f,est}$$

$C_{f,est}$ is the prior estimate for completing the feature.

We can substitute this in our equation as follows:

$$\begin{aligned}\Delta E(V) &= O - E(R2) - E(C2) - (O - E(R1) - E(C1)) \\ &= - E(R2) - (E(C1) - C_{f,est}) + E(R1) + E(C1) \\ &= - E(R2) + C_{f,est} + E(R1) \\ &= - \Delta E(R) + C_{f,est}\end{aligned}$$

In other words, once a feature has been completed and integrated, its value to you is equal to the reduction in risk plus the original cost estimate for the feature.

If the feature is completed *exactly on-budget*, then the original cost estimate was accurate, and then – and only then – can actual cost be used in lieu of estimated cost as a component of the measure of a feature's value.

Adjusting For Concurrent Progress On Other Features

In real software systems it is usually more complicated than this. In particular, it is usually the case that some features cannot be tested unless other features are also built and in a testable state. Interdependence of features is illustrated in Figure 8 on page 126.

This fact of life complicates value measurement because it means that feature values are not independent. To deal with this, one can refrain from attempting to estimate the individual value of such closely related features and instead estimate the value of the entire set of features that must be built and tested together. The entire feature set then has a value, but the value of each feature in the set has an undeterminable value. One should not attempt to measure the value of such granular features anyway: time would be better spent on helping the team to actually build the features and verify them.

Liquidation Value

It might occur to you that if precursor features have no independent business value, then why is it that one often can sell the code of an incomplete system? I myself have seen this done. When an incomplete, non-operational system is sold, it is *liquidated*.

When one buys a system or some of its components in a liquidation sale, the value that one receives depends on the intended application of those components. The components obtained reduce the risk and cost of building the business capability that the buyer envisions. Therefore, when I say that a precursor feature has no *independent* business value, what I mean is that it has no operational value by itself. That does not mean that it does not have market value (economic value), and that value is based on the extent to which it reduces the risk and cost of completing a capability that the buyer needs.

It follows then that the liquidity of a feature has a lot to do with its value. That is, if a feature's code can be sold independently, then its value can be estimated fairly accurately, based on what people are willing to pay. Features with minimal operational business value tend to be fairly non-liquid and sell for a low price – if at all – while those with clear independent business value tend to be saleable and therefore (by definition) liquid. The price and liquidity difference reflect the difference in risk with respect to the actual value of the feature. Thus, one can view the distinction between a demonstrable capability (a completed "Minimum Marketable Feature") and a feature with no independent operational value as merely a difference of liquidity.

Modeling the Value of Continuity: Making the Case to See a Strategy Through

IT is famously beholden to organizational changes that occur outside of it or above it and that lead to sweeping changes in IT strategy – without the input of IT. Consider for example, the hiring of a new CIO, the decision to cancel a project that has just begun to bear fruit, and a merger with another company and consequent consolidation of IT activities: all of these kinds of changes usually have a massive impact on IT's ability to realize value on its prior investments, yet it usually has little say about these "strategic" decisions, even though the impact of IT affects the organization's ability to realize value.

There is an opportunity cost to abandoning strategies before they have completed. This cost is the expected value of the partially completed capabilities, minus the expected value of those partial capabilities when applied to the new strategy of the organization. That is, it is the liquidation value in the context of the new strategy.

Firms usually see unfinished investments as sunk costs and therefore irretrievable. To confuse things more, the accounting value of unfinished business capabilities is zero, because those partial capabilities are not usually assigned a market resale value, and since they are not generating revenue, they have no tangible value. Executives therefore often conclude that there is no harm in "starting over" at any point, when in fact to abandon a strategy is to discard the economic value of the partial capabilities.

Of course, this might make sense if there is reason to believe that the capabilities will not actually realize value: that is, if the expected value of the investment has dropped to the point where it is no longer a worthwhile investment, even given the work that has been accomplished.

Another factor that often confuses matters is a failure to value the investment in human knowledge. I have seem hundred million dollars projects cancelled because executives lost confidence in the cost and benefit balance of an effort that was finally starting to succeed: just as staff had got traction in delivering true capabilities, the plug was pulled because a strict analysis, at that point, of the cost and benefit showed an imbalance. However, my own assessment was that substantial value would have been delivered – a value well worth the investment – if management had placed a value on the knowledge and consensus of approach that had been developed among the hundreds of staff.

It is critical that IT have concrete arguments when trying to make the case to stay the course. Often the drivers are political and therefore very difficult to challenge. However, it is in the organization's best interests to present an honest case that assesses the value of each option, and that includes the lost opportunity costs of abandoning projects.

Another Example

Suppose a company called Gigantic E-Commerce[1] wishes to add a capability to allow business partners to market their own products through Gigantic's website. The benefit to Gigantic is that it will receive a percentage of each sale. Gigantic estimates the expected business opportunity to be $25M over three years.

In order to make this capability available to partners and customers several features must be added to Gigantic's systems and business processes.

Gigantic divides the overall capability into two distinct feature sets: (A) the Partner Interface component set; (B) the customer interface enhancements, and (C) modifications to the back-end database-related components. Each of these feature sets involves a substantial amount of work, but neither provides business value on its own: value is only realized when both features sets have been integrated and deployed as a new business capability.

Gigantic decides to use an agile, incremental development process in order to design and implement this new capability. Toward that end, Gigantic needs to develop a model for estimating a reasonable theoretical business value for each feature as it is developed. Otherwise it has no way to assess progress. Gigantic has learned that assessing progress based on which tasks have been completed does not provide sufficient incentive for teams to be inventive with respect to how they accomplish goals. Project plans must be fluid and adaptable, especially since Gigantic expects to work with prospective partners who might request for modifications to the features as they are being built. Therefore, a value-centric planning approach is critical to assessing progress toward the end goal.

Gigantic decides to first implement feature sets A and B, and then use these to solicit detailed feedback from customer service and from partners. This feedback will perhaps be used to fine-tune the feature sets. During this period work will proceed in parallel with feature set C, but feature set C will likely not complete until after feature sets A and B have been completed.

At the beginning of work, Gigantic estimates that feature set A will require about $1M to complete, and feature set B will require about

[1] This company name is fictitious. Any similarity between this name and the name of any actual company is purely coincidental.

$2M to complete. These are shown in the "Prior Expected Cost To Complete Feature" column in Table 15.

Gigantic also knows based on past experience that these efforts sometimes go nowhere, either because objectives change or because of technology problems. It therefore estimates that there is a 10% chance that feature set A will not succeed, causing the entire initiative to fail and eliminating the business opportunity. This 10% translates into a $2.5M risk for the opportunity, as well $1M for work on the feature set which will have been wasted. In fact, all prior work on the capability would have been wasted, but that risk will cancel out since we are estimating incremental risk for each feature.

A similar analysis for feature set B leads Gigantic to conclude that there is a 10% risk for feature set B, translating into another $2.5M plus $2M in costs. These risks are shown in the column "Prior Expected Risk To Complete Capability" in Table 15.

Gigantic then completes work on features sets A and B, with work proceeding along on feature set C. Gigantic is fortunate that partners have provided feedback indicating that no changes are needed to these feature sets, and so the work on them is considered to be completed. Gigantic's progress can be summarized by the column "Value Of Feature" in Table 15: the risk associated with each feature has been eliminated because they have been completed successfully, and the value of each feature set can be estimated by summing the risk that has been mitigated and the prior expected cost. The total value of feature sets A and B are shown in the last column and are seen to be $11M, and this value results from the reduction in risk to the opportunity as well as the value of the work done.

The actual build costs can be used to compute a theoretical ROI for each feature, based on the feature values. The build cost should include the cost to integrate and fully test the feature set.

Table 15: Example of estimation of theoretical feature value

Features Completed This Past Cycle	Prior Expected Risk To Complete Capability	Prior Expected Cost To Complete Feature	Actual Build Cost Of Feature	Value Of Feature
Feature Set A	3.5M	1M	.8M	4.5M
Feature Set B	4.5M	2M	2.4M	6.5M

Combined Feature Sets	8M	3M	3.2M	11M

It is important to understand that these feature values are only theoretical. In reality, a feature has no value by itself: only capabilities have value – by my definitions. A feature should not be treated as a tangible asset, because it is unproven. Capabilities are tangible assets: they have measurable business value. A feature is by itself an unrealized investment. If you need to estimate a liquidation value for the feature, as is the case when a company sells the IP assets for a product that has not been completed, the feature should be packaged and proven sufficiently to establish a tangible value for it.

The values calculated here for features are only for the purpose of tracking incremental progress in a manner that is as realistic as possible.

Relationship To Earned Value

The Earned Value Management System[1] (EVMS) is mandated for many US Federal Government agencies, according to US Office of Management and Budget (OMB) policy – the same entity that oversees the Federal Enterprise Architecture (FEA) standards. The National Institutes of Standards and Technologies (NIST) provides a EVMS standard that can be purchased online.

EVMS relies heavily on the ability to define objective metrics for assessing progress. However, EVMS determines value based on the planned cost. In essence, the value of a component is determined by the prior estimated cost of work required to meet certain measurable milestones such as demonstrating that components have passed their acceptance tests. This is a poor way to measure value, because it cost bears little relationship to business value, especially if one is trying to measure the value of features that do not have their own independent business value. In fact, very often the most expensive components are those that add the least business value. Instead, one must take account of the risk to the value of the ultimate planned business capability. Only then can one say that effort is valued in accordance with expected business value. EVMS is more about holding contractors on large projects to the plan rather than about measuring actual value, and that is why EVMS is used for large government projects. The business sector

[1] EVMS is defined by ANSI standard ANSI/EIA-748-A-1998 (R2002). A tutorial can be found at https://acc.dau.mil/evm.

requires more focus on actual value, and the ability to modify the plan midstream is much more important than holding a vendor to a plan.

Relationship To Total Value of Ownership

In recent years the concept of Total Value of Ownership (TVO) has taken root.[1] The TVO idea is designed to contrast with the concept of Total Cost of Ownership (TCO) which focuses on the complete lifecycle cost of IT systems (or any systems for that matter). Since a failure to realize value can be seen as an economic cost, the shortcomings with TCO are more in how TCO has been applied than in the concept itself. Thus, TVO does not really add anything new except for a renewed emphasis on business value.

TVO is new enough that one cannot say much about how it has been implemented, but early reports seem to focus on the impact of IT systems on revenue. This is an important element, but one also needs to consider the impact of decisions on risk, and the impact that risk mitigation has on expected value.

Comparing Alternative Implementation Strategies

IT decisions often involve alternatives that boil down to "build it the 'right' way" or "build it the cheapest way". The most sensible choice depends on the balanced business impact over the timeframes that are of concern. However, the business impact of IT improvements can often be boiled down to lifecycle cost – assuming that one truly includes the full lifecycle for the IT system within the organization and the impact (perhaps many years later) on other initiatives during that timeframe. In other words, one should consider the impact on future agility and evaluate the importance of that with respect to other more immediate concerns.

To do this analysis one must define the alternatives, and then one must consider the various hypothetical future scenarios that can be expected to evolve and how costs and agility would be affected. As a simple example, consider an organization that has two systems called A and B, as shown in Figure 18.

[1] Ref. [Luftman05].

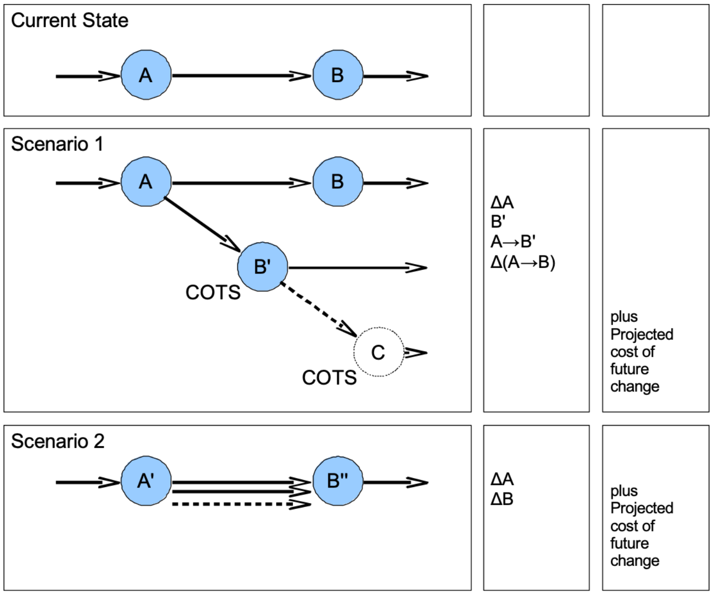

Figure 18: Example of computing lifecycle cost for alternatives.

In this example, a business unit in the organization would like to add a "commercial off-the-shelf" (COTS) system B' as shown in scenario 1. The COTS system would allow them to handle a new class of business, thereby tapping a new market. This represents a substantial opportunity for the organization in many tangible and intangible ways.

The technology-savvy members of the organization point out that the logical function of the COTS system is the same as the existing system B. The difference is that the COTS system incorporates some business rules that the existing System B does not. However they acknowledge that it would take awhile to enhance B to incorporate those new business rules, and that buying the COTS system might be a shorter path to market.

The dilemma is that adding the COTS system complicates the architecture of the enterprise: the organization would then have two systems that do similar things, and each with a very different architecture and methods of maintenance and integration. Further, if new rules must be added to B in the future, it is likely that those same rules will have to be added somehow to the COTS system, perhaps by asking the vendor of the system to add them or by adding a pre-processing or post-processing system in front of or after the COTS system. The technology-savvy folks are very concerned about the prospect of growing complexity.

The business side does not understand these arguments, because "complex" does not mean anything in business terms. So what if it is complex? Isn't everything complex anyway? The tech-savvy folks respond, "but complex means greater lifecycle cost." The business folks respond to that by saying, "Show me: otherwise I cannot evaluate that in terms that I understand; and further, give me an alternative, since not going after the new business is not an option."

The tech-savvy folks then have two tasks before them: (1) to devise a reasonable alternative, and to (2) show that the alternative is superior when one includes the combined impact of future business and total future lifecycle cost impact, *including the expected impact on the ability to capitalize on future business opportunities* (for example future ideas for new products or partnerships) that have yet to reveal themselves but that will surely come along.

This is a tall order, but the tech-savvy folks work in a very progressive capability-focused organization that has developed techniques for modeling these things, and they are respected by the business for that reason; so they devise scenario 2 as shown in Figure 18. Further, they hypothesize that once either scenario has been put in place, that there will be a new, as-yet unforeseen business opportunity, and that the impact for scenario 1 will be to require yet another COTS system C to be installed in parallel to the earlier one (B'). The impact for scenario 2, which assumes a strategy of enhancing System B, will be to enhance B again, as shown.

Thus, in Figure 18 the dotted lines represent hypothetical unforeseen business needs, and the way that the need is addressed in scenario 1 is different from the way that it is addressed in scenario 2 because each scenario assumes a different strategy: scenario 1 assumes a strategy of adding specialized systems, whereas scenario 2 assumes a strategy of extending and refactoring existing systems. These scenarios can be compared in general cost terms. **The hypothetical future scenario with the lower cost indicates the more agile strategy, since cost translates into effort**. The value of the enhanced agility can be taken to be the anticipated value of the hypothetical future opportunity, reduced by the cost and risk of realizing it. As we know, increased agility means reduced risk, time to respond, and cost for capitalizing on new opportunities.

This is a very simple example – a trivial one at that. Real business computing environments usually have a multitude of interconnected systems, and the interconnections represent enormous investments; more interconnections in the future represent even more investment and disentangling those interconnections can be an effort that escalates geometrically with the number of interconnections. To truly compare future scenarios one must define a model scenario that is representative of the systems and kinds of interconnections that are in place. This process is not too different from the practice of threat analysis employed in the security profession. The situation there is that instead of threats, we want to identify possible future events and outcomes.

If one is ambitious, one can even build a simulation model of these outcomes, estimate the probabilities of various events, and then run the simulation to see how the events combine over time to yield various overall outcomes. The frequencies of these simulated outcomes are in effect probabilities, and these probabilities can be used to estimate the expected value of alternative strategies today, since the choice of strategy will influence the outcome. This sophisticated approach is used routinely by the investment community. For enterprise architecture analysis, it is usually sufficient to merely identify the relevant future events, such as new product introduction and market changes, and project future technical outcomes by hand.

Relationship To Real Options Analysis

The Real Options Analysis (ROA) financial valuation methodology, mentioned briefly previously, is a class of analytical methods for comparing alternatives that have uncertain outcomes.[1] The ROA approach applies investment present value principles in a probabilistic context.[2] The fact that IT investments can be modified over time is taken into account, generally by allowing for a set of "options" such as (a) abandoning the original investment, (b) increasing the rate of investment, and (c) deferring a decision until events indicate which course of action is better.

ROA originates from the concept of extending the techniques of financial option analysis to other domain. These techniques are

---

[1] For general information on the Real Options methodology, see [DixPin94] and [Copeland03].

[2] To an economist this statement may sound contradictory, since Real Options Analysis (ROA) is viewed as a generalized method of which Net Present Value (NPV) is a special case; however, the phrasing here is meant to introduce the concept to readers who do not have a background in financial analysis.

based on the famous Black-Scholes model and formula. Financial options are purchasable contracts that allow the buyer to either buy (a "call" option) or sell (a "put" option) a financial instrument for a pre-agreed price (the "strike" price) at some point in the future. However, when applying the option purchase analogy to IT, it is as if one had a financial option for which the stike price is uncertain and the option price is uncertain, since with IT investments one does not know with certainty either the eventual tangible value that will be realized from the investment, nor does one know the investment cost with certainty until the project has completed. These uncertainties make the strict application of options analysis somewhat limited, and in practice more general approaches are needed that account for the many uncertainties. However, the very idea of treating investment as a probabilistic endeavor was a breakthrough in thinking.

As of this writing the ROA approach is not yet widely used in IT, but it is used extensively in businesses that have very large capital investment decisions such as the oil drilling, biotechnology, and telecommunications. However, some have attempted to apply it to IT [1], and its use is definitely on the rise.

One of the challenges in applying a Real Options approach is that one must estimate the opportunity value of various courses of action: that is, one must estimate the probabilities of future outcomes and their up-side and downside values. This is the problem that we ran into above when trying to put a number on the value of unforeseen future opportunities. Thus, that is the central challenge. In some circumstances it is worthwhile to utilize stochastic simulation of future events to see which general outcomes are the most likely. This approach is widely used in the securities industry to project the value of different investments given that interest rates and other indicators might follow various paths of varying likelihood; but in the end it is still an *educated* guess. In the next section I discuss stochastic analysis and its applicability to IT strategy value analysis in more detail.

Within the context of a software development initiative, some of the options include continuing or canceling projects, as well as intervening in projects to improve (hopefully) their progress. In the context of a single project, options include how to implement the software: that is, architectural decisions. These decisions impact the

[1] Ref. [Kazman01].

future capabilities, agility, and cost structure of the organization. Therefore, future events are important determinants of whether the choices made turn out to be good ones. The kind of scenario analysis described in the prior section is a useful tool for considering what the possibilities might be. In order to evaluate the expected risk and opportunity values of these possible future outcomes for a given architectural choice, one must determine probabilities for these outcomes. This is the essence of Real Options Analysis. Statistical simulation, known as "stochastic analysis", is a tool that is used in many industries for this kind of problem, but as a starting point one at least needs to define the future scenarios. As the organization becomes more experienced in applying this kind of analysis with IT projects, it can start to consider simulation and the more advanced methods used for ROA.

The future cost of abandoning work is not merely the lost investment: it includes the costs of *disentangling* what was built.

When it comes to IT it is important to consider that when one chooses an option to build an expedient solution that violates an architecture, *the future cost of abandoning that decision is not merely the lost investment: it must include the costs of disentangling what was built*. For systems that have been in use for awhile and which have been interconnected to other systems, this cost can be very substantial. See the discussion in the section beginning on page 372. Thus, the cost of abandoning a project may extend beyond the sunk costs and lost opportunity value.

Another important lesson from Real Options Analysis is that there is inherent value in the flexibility to postpone a decision. That is, one can increase one's chance of success by postponing certain decisions (e.g., whether to abandon a project), and this value is tangible. The rationale is that if one waits to decide, one will have more information available and therefore make a better decision. The flip side is that there is a cost also in postponing certain decisions, in particular if work is proceeding in a direction that might be the wrong direction. Modeling these competing risks is the challenge.

Stochastic Analysis

Stochastic analysis is becoming more prevalent as a tool for estimating the business value of decisions and strategies. The basic approach is simple: try to conceive of everything relevant that might happen that will impact your business plan and strategy, and model those events as event trees; then estimate probability distributions for the events, and simulate the scenario a large number of times. The distribution of outcomes that result from the simulation are then used as the basis for computing the value of the strategy. You can repeat this process for various alternative strategies, and then compare the outcome distributions that result.

The event modeling portion of this approach is somewhat analogous to security threat modeling. In threat modeling you try to envision every way that someone might attack a system, and estimate the approaches that are most likely to succeed. Threat modeling tends to stop there, because the purpose of threat modeling is to identify weaknesses and then fix them. However, in some cases it is not possible to fix a weakness right away, and so it is necessary to evaluate the risk posed by the weakness, and the cost-benefit of mitigating the risk. Simulation can be used to model that, if one assumes a probability distribution for the success of an attack, for a given mitigation strategy.

Let's consider an example. Figure 19 illustrates a simple event model for security risk mitigation.

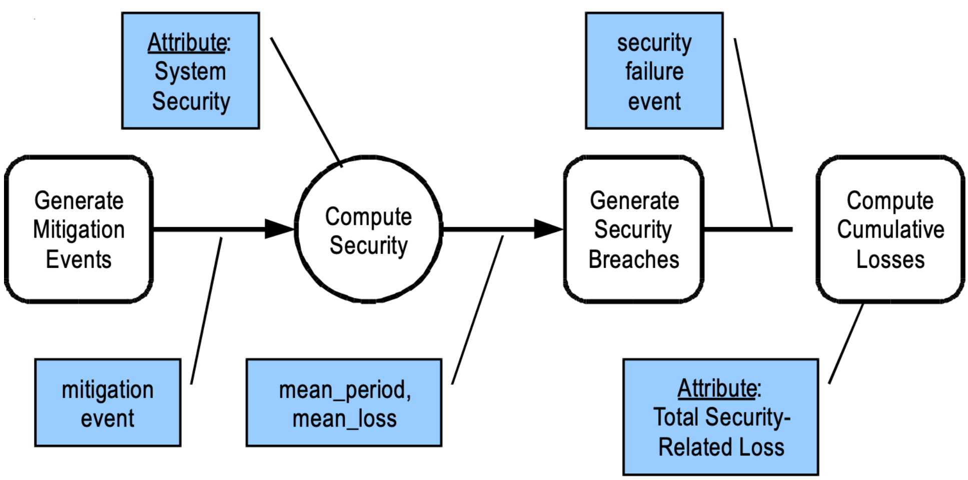

Figure 19: A simple event model for security risk mitigation.

The purpose of the model in Figure 19 is to identify the things that can happen in the lifecycle of a system, for which there is known to be a security risk. It is assumed that the organization wishes to evaluate the return from different levels of investment in risk mitigation, such as better firewalls, more analysis of application code, etc.

The first rounded box, "Generate Mitigation Events", represents the deployment of practices that reduce security risk. For example, if the organization deploys a new application-level firewall, this might greatly reduce the risk of an entire class of threats. Such a deployment would be an event in the model. The "Generate Mitigation Events" box produces such events according to a distribution over time, and each event carries with it parameters that characterize the impact on overall security.

The bubble "Compute Security" is merely a function that reads the mitigation events and calculates the net security level for the organization, according to some scale. This security level is then input to the box "Generate Security Breaches", which uses the security level as a parameter in its distribution of events that it generates. Each of the events generated by the "Generate Security Breaches" box represents an actual security failure, such as a successful hack, or a loss of data due to theft.

The events produced by the "Generate Security Breaches" box are received by the "Compute Cumulative Losses" box, which tallies all of the security failures that occur, along with the value of loss, to produce a total loss over a simulated period of time. This total loss can be compared with the size of investment in the mitigations. Further, by tweaking the parameters of the "Generate Mitigation Events" box, one can compare various mitigation strategies, such as whether investing in firewalls is better than investing in application threat analysis. Of course, the old adage "garbage in, garbage out" applies, and the model is no better than the assumptions embedded in each box; but at least the assumptions are explicit and can be examined and discussed, and the results compared with experience, until confidence in the model is achieved.

There is also the complexity of selecting probability distributions for the event generators (the rounded boxes). The distributions can often be estimated based on historical information. For example, if one needs to estimate the magnitude of impact on security as a result of using a certain type of mitigation, there are industry numbers for the effectiveness of different kinds of security, and the organization's historical rate of security attacks can be used to create a weighted expression for the projected impact. A great deal of judgment is required here, and one should not expect to get it right the first time. It is as with marketing: one must make educated guesses, and businesses are well-accustomed to that type of uncertainty.

It also turns out that the exact shape of a probability distribution seldom matters much in the final results: if it does, this might be a red flag that the model has some important edge conditions in it. In any case, the most general distribution for this purpose is the gamma distribution, which can be used to model a variety of stochastic event processes. In this book I am not going to go into the details of distributions: most simulation tools provide such distributions out of the box and provide guidance on selecting them.

The example shown here is extremely simple to illustrate the technique – too simple in fact to make the case that this technique is worthwhile. In fact, stochastic modeling is extremely powerful and is worthwhile in many IT decision-making situations. It provides the ability to explicitly model the interactions among decisions and the holistic impact on value creation (as well as cost), while integrating both risk and opportunity. Further, once the model has been created, it provides a basis for discussion, in terms of the model parameters and interactions. Over time, through trial and error, a consensus can be developed about the validity of the model.

Stochastic modeling can also be used for generating input to portfolio analysis. In the paper [Bhadra06], dependencies between IT portfolio investments are represented by a statistical correlation value. Correlation values (as "correlation coefficients", or as "co-variance") can be determined by simulating scenarios in which inter-related projects proceed with varying levels of investment. The correlation coefficients can then be used in an overall enterprise portfolio model. This is beyond the scope of this book, but I mention it to show the relationship to portfolio analysis.

The *Expressway*™ Simulator

There are commercial simulation tools available that can be used to create stochastic models. However, they are not generally targeted at IT architects, and so the modeling paradigm tends to be mismatched to what architects expect. For example, architects like to work in terms of structural models, and are accustomed to programming concepts. In contrast, most of the financial modeling tools available are based on a spreadsheet paradigm, which is generally less familiar to architects. Further, the most popular tools are proprietary.

In order to address these problems, I created a tool called *Expressway*™, which can be downloaded for free from this book's website. While there is a commercial version (which is necessary to fund development of the free version!), the free version is fully usable: it merely lacks many of the features that are useful for collaboration and long-term work across an enterprise.

This is not a book about *Expressway*™, so I am not going to provide a tutorial here about the tool, but there is ample documentation on the website. Further, *Expressway*™ has pre-defined generic (and customizable) implementations of all of the architectural value models described in this book.

More Value Models

So far in this book I have provided guidance for how to model the value of things such as work on incomplete projects, risk mitigation efforts, anticipated business opportunities, flexibility in terms of its contribution to business agility, and others.

IT architects worry about many things that have strong business relevance, but that are commonly expressed in IT terminology that is meaningless to business. These include things such as decoupling, component roles, constraints about what components are allowed to do, constraints regarding communication among components, decentralization versus consolidation, generality of function versus specialization, and refactoring. (The *Expressway*™ tool has generic business value models for all of these.)

There are also some technical issues which business is familiar with, because these issues have bubbled up to the business level and are easily understood: these include data quality, and the ability to make better decisions because of more transparency with regard to business rules and correlations within data – this has to do with "business intelligence", and it hinges on the ability to correlate data, which often must be built into applications because cause-and-effect correlations are impossible if they are not captured proactively.

IT is in the position of either (a) keeping its decisions about these issues to itself, or (b) explaining these decisions to business. Common practice is to do the former, except that the decisions often have a large impact: e.g., a decision about whether to consolidate applications across geographically distributed operations to save money, at the expense of the operational flexibility that those geographic sites are accustomed to. Which is the better choice? If one does a cost analysis, consolidation will always win; but what it the business cost? What is the value of the lost opportunity to be more flexible, and to be able to respond more quickly to changing business needs? To answer this, IT must help to figure out the value components, because many of them have to do with how IT systems are planned and operated: that is, IT value is not just about *what* to do, but *how* to do it.

Let's take "decoupling" as an example. Decoupling is the buzzword that one hears the most often when it comes to justifying decisions to generalize components. What does it really mean, and what is its actual value?

Decoupling is typically justified as a path to being able to reuse components and to untangle behavior. It is an expression of the degree of separation of concerns between two components. It has implications for reusability, reliability, testability, and maintainability. Therefore, its benefits should be expressed in terms of these things.

On the other hand, there is a cost associated with decoupling components, in terms of performance, effort, and complexity. The impact on complexity can be positive or negative: if two components are highly inter-twined, maintenance might be difficult because the tangle is confusing; but if too much separation is created by defining more types of components than are really needed, complexity can go up as well. The right balance is determined by judgment.

That judgment needs to be applied to estimate the impact on business processes pertaining to future development, maintenance, and testing. One should also assess the likelihood of actual reuse: it does no good to make a component reusable if it will not actually be reused.

Discussions with business should be in terms of these tangible qualities: costs of maintenance, savings from reuse, increased availability due to better reliability, and so on. Business should not be left out of these tradeoffs.

Many of these kinds of decisions are made at multiple levels. For example, one might be considering refactoring a single object class, or an entire system interface. Business does not need to be involved in the low-level decisions, but they should be involved in the strategies by which these decisions are made: does the *system* need to promote reuse? Does the *system* need to be highly reliable? Does the *system* need to be built to be maintainable? Once decided, based on a value projection, IT can execute these preferences, and attempt to live up to the projections.

The Value of Better Decisions and Data

It is difficult to justify the value of better decision-making ability. Access to higher quality information and to better models improves decision-making. The value of this depends on the decisions that remain to be made.

Over time, there will be an impact that is hard to predict, because after all, you don't know what you don't know. Therefore, the value of better decision-making capabilities is best justified by trying it on a small scale first. A pilot project to try new decision-making capabilities, through better data quality, better data analysis, and better predictive models is often the way to go. If this is not possible, one can look at other parts of the organization where better decision-making capabilities were established and assess what the impact was. If this kind of comparison is not possible, management might need to accept (or reject) the projected benefit based on instinct.

Value and Risk Management Practice Frameworks

As we have seen, risk is merely negative opportunity, and so risk mitigation is an investment in reducing the expected value of a negative opportunity. This allows us to model risk as any other opportunity and to use investment analysis to compare it with other uses of resources. Further, any business opportunity – positive or negative – is but one of any organization's many concerns that must be balanced to maximize shareholder value or other goals over a planning horizon.

In other words, every single action by an organization should theoretically be evaluated in the context of a unified model that integrates opportunity with risk, and balances those investments with other investments that are made by the firm.

This sounds logical, but of course it is cumbersome and infeasible to treat every action as an investment. For some actions it is just not worth the trouble. To make the distinction I will refer to "practice-level" actions as those activities that normally fall below the level of planning in question, and for which an analytical value analysis is not worthwhile.

For example, if one is planning a program for introducing a new product line, the plan might include a series of initiatives in different parts of the organization. In such a plan, details about the operation of those initiatives are below the radar screen and would be considered to be "practice-level". However, to those planning the initiatives, such details are very pertinent, and would be considered to be plan-level. On the other hand, some things are so granular that they are inherently practice-level. For example, the manner in which software is deployed should be treated as practice-level unless the deployment is for a strategic system and is unique in some way.

So why does this distinction matter? The reason is that it is inflexible and cost-prohibitive to perform a cost-benefit analysis for small decisions, and yet guidance or rules of thumb are necessary so that good decisions can still be made by less experienced staff. For example, suppose one has to deploy a software system, and one must choose what methods to use, depending on the value at risk represented by the system. Many organizations use categories for different levels of risk, such as "Tier 1", "Tier 2", and so on. The idea is that a "Tier 1" system is high risk or mission-critical, and therefore it should be deployed and operated in a very well-controlled manner. In that case, a "Tier 1" system would have an associated procedure for deployment that is very comprehensive and well-defined, whereas a "Tier 2" system would have a less comprehensive procedure. That is, the *practices* to be used for deploying a "Tier 1" system would be different – more strict – than those for a "Tier 2" system.

The question then is how does one select the right types of practice-level actions? What should the criteria be for selecting, for example, the right amount of software testing? The right types of use cases? These are all detail-level issues, but they are very important in the aggregate for the success of the organization when the practices are reused across many systems, and these low-level decisions must be based on business value.

To address this question an organization needs a set of prior decisions on all of the major practice issues, cast as a table and driven by criteria that practitioners can easily decide. For example, a set of criteria can be used to determine if a capability is "high-risk" from a security perspective, and all such high-risk systems might require a predetermined level of testing. The analysis can be done once: a model of a hypothetical high risk system can be defined, and the appropriate level of mitigation determined for each of several categories of system value based on business value analysis. This provides a benchmark for the level and types of mitigation that are required for all similar high-risk systems. The criteria for the high-risk system can then be reused to identify which other systems are high risk, and the same mitigations can be applied without having to redo the analysis.

The result is a decision table for which practices to use that can be used broadly to make quick decisions that are based on an integrated model of risk and business opportunity value. This type of table is an invaluable tool for guiding on-the-ground actions. It provides a kind of maturity framework – but in a much more flexible manner since it is sensitive to need and can be updated whenever the analysis model is updated.

To reiterate, the in-the-field usage of such a tool is then a two-step process: first use characteristics of a business process to categorize it, and then look on the decision table's list of mitigation practices for all practices that fall within those categories.

The details of how to prepare such a table is beyond the scope of this book, but it is fairly straightforward analysis that is little different than the techniques provided in this chapter. Those in the insurance and credit industries will be familiar with this type of analysis. The important point here is that value-based decision-making is possible at all scales of activity, and that at granular task levels it can be implemented via practice frameworks and decision criteria for choosing practices. Note that I do not use the term "best practice", because that term implies that one size fits all, when in fact what is "best" for one situation is not necessarily is not necessarily best. In fact, it is not about using the best practice: it is about using the *right* practice.

Risk management frameworks such as COBIT [1] provide very useful catalog of types of risk. For example, COBIT 4 defines a practice called "DS5: Ensure Systems Security". The specific practices that should be employed to implement this practice category should be

[1] Ref. [COBIT4].

very much more stringent for a high-risk system than for a low-risk system. Therefore, judgment or value analysis is required in applying any maturity framework.

Maturity frameworks such as CMMI [1] have the philosophy that an organization should basically "know what it is doing" by virtue of having defined its own processes. A decision-making practice table is certainly aligned with such an approach, and the use and maintenance of such a tool is consistent with the CMMI Maturity Level 5 category in which the organization is actively tuning and optimizing its processes.

Measuring the Value of Knowledge

Knowledge is something that exists only in the heads of people – at least as of this writing. Knowledge is much more than information: knowledge implies a level of conversance about a subject, and therefore it also implies an understanding of the subject. Knowledge does not need to be learned: it is the result of learning. That is, one reads information in the form of documents, views information in the form of presentations, and hears information in the form of explanations, and as a result one builds knowledge.

This is why knowledge has so much more value and potency than mere information: knowledge is ready for action. The potency for action is what gives knowledge its value in business. Therefore, the value of knowledge depends on the potential for action: knowledge is only actionable in business if it pertains to opportunities for value creation or risk reduction.

The business value of knowledge is the value of the opportunity that it represents for action.

The business value of knowledge is the value of the opportunity that it represents for action. Rather than try to measure the value of knowledge based on the cost of acquiring it, we should measure the value of knowledge based on the opportunities that it represents. However, we must account for the fact that knowledge is usually not sufficient to perform a task: one usually needs other resources, and so the value of the knowledge is not independent of those resources. The value of that knowledge can therefore be estimated

[1] Ref. [CMMI].

in a manner similar to the way in which we have measured the value of features that have no independent business value (see the section beginning on page 245). That is, the value of the knowledge is the *amount by which it increases the expected value of what the knowledge can be used to accomplish*.

Another aspect of knowledge is that its integrity changes over time. Over time people generally lose detailed information. However, they can also develop a more mature and holistic understanding of a subject if they have remained involved in the subject in some manner and have therefore learned more about related subjects so that they can make mental connections that foster a broader understanding. Thus, over time, knowledge can be maintained and honed, it can be matured, or it can be lost. Which happens depends on the activities of the person in the time since the knowledge was acquired.

It would probably not be effective to try to model the loss or maintenance of knowledge, at least for our purposes here. To do so would be to create a very complex and theoretical model that would need to be calibrated and that would be subject to inaccuracy in individual cases. It is better to take a more practical approach of using subjective evaluation. Specifically, if one asks, "How effective will person A be at task Y the day after they complete related task X?" then we have a basis for assessing the impact of allowing person A to experience a period of delay between tasks during which they work on something unrelated. For example, if a year goes by and then the person turns to task Y, will they require a ramp-up time to refresh their knowledge? Similarly, what is the ramp-up time for someone who is new to the task, even if they have generic knowledge about the subject domain of the task?

These are important and relevant questions that managers ask themselves all the time, but we need to assign value to knowledge so that it is accounted for as an asset that is accumulated as a side effect of projects. In fact, such knowledge is not a side effect, but a direct result of efforts by mentors to instill knowledge in the minds of those who they mentor.

I recommend tracking knowledge as an asset. That is what it is, and if we do not keep track of where it exists and what value it has, we will not properly manage it. Knowledge of systems or processes represents an important component of the capabilities provided by those systems or processes. It should be inventories just as we inventory our business processes and our systems. That is, it is a kind of "meta" data – data about data, processes, or their elements – that one should track and associate with each individual in one's staff. Important attributes of that meta data include the nature of the knowledge and how current or robust it is. Once we track this, we are in a position to monitor it and ensure that it remains at a level that is needed to maintain business processes.

To evaluate a knowledge asset, one must assume the theater of action for the knowledge: how will the knowledge be used? Will it be used to operate an existing business process? To maintain a business system? To enable future changes to an existing business system? To enable entirely new but related business opportunities? To mitigate risk with regard to existing processes or systems? This means that the value of knowledge to us is not constant: it changes as opportunities presented to us change.

For each opportunity scenario in which the knowledge might be used, one can evaluate the amount by which the possession of the knowledge increases the expected value of an investment in pursuing the opportunity. This applies whether the opportunity is an opportunity to increase revenue or an opportunity to mitigate risk and therefore reduce expected loss.

Consider for example a situation in which a project to create a new business system completes and the team is disbanded. The members of that team bear knowledge about the business opportunities associated with the project. As they go their separate ways, they retain that knowledge for a period. If an opportunity appears to add new capabilities to the business system, one can assemble an entirely new team consisting of generic skills, or one can attempt to obtain the original team. It is obvious that the original team will have a shorter learning curve and therefore reduce the cost of completing the effort. However, what is often not considered is that the original team might reduce the risk of failure. Therefore, their presence increases the expected value of the investment. The chance of failure adjusts the expected return. That value is higher for if the original team is used. The difference represents the unique value of that team.

Dealing with Poorly Articulated Financial Goals and Strategies

Any business value metric must ultimately roll up to the core goals or strategies of the organization. For a profit-focused incorporated entity, this is always shareholder value. However, the time period over which value is maximized depends on the planning horizon of the organization. For a non-profit or governmental entity, business value metrics must be in terms of success with respect to the mission of the organization.

The important thing is that the measurement approach and metric have consensus and be the result of an investment of time, effort, and reputation on the part of both IT and the major operating units. Otherwise they will have no staying power when difficult decisions arise.

Consider the following example core strategies for a for-profit public corporation:

1. Introduce additional competitive products.
2. Manage market risk better.
3. Reduce financial non-compliance risk.
4. Reduce information security risk.

These strategies are undoubtedly rooted in shareholder value, but the board and executive committee has analyzed the company's situation and decided that these four strategies will maximize the expected shareholder value, taking account of the impact on shareholder value of these risks and opportunities. Therefore, for the purpose of execution planning, these four strategies are the goals that management must work toward.

The problem right off the bat is that there is no guidance in the four strategies for how to trade off each of them against the others. For example, what should one do when a decision has two choices, and one choice maximizes the chance to introduce new products, and the other greatly reduces market risk? What is the value of the new products when compared with the reduced market risk?

Rather than try to crack that nut right off the bat, it is best to start by trying to define a metric or scale for each of the four goals. For example, goal 1 can be expressed in terms of the expected value of an additional product, as shown in Table 16.

Table 16: A metric for measuring the value of the introduction of new products.

Goal	Metric
Introduce additional competitive products.	Expected value of product over ten years, adjusted for direct and indirect costs.

This metric might be hard to compute, at least initially, but it defines the measurement and thereby gives planners a starting point. Also, many organizations do compute this kind of metric routinely for their products, but what is often left out are the many indirect costs. By defining this metric for IT, IT now has a chance to identify the indirect costs and tradeoffs that they know about.

Let's consider goal 3 now. I will skip goals 2 and 4 because they are somewhat redundant in character with regard to 1 and 3. Goal 3 has to do with the risks associated with failure of a financial audit, as well as any costs associated with financial regulation compliance. This is a large cost for a public company because the rules are very complex and there is a very complicated responsibility for operations to report their business transactions, which the accounting department must then translate into official ledger entries. Note that these costs are included with the risk because a failure to pass an audit results in costs for remediation of the failure.

The value of financial compliance can be expressed as shown in Table 17.

Table 17: A metric for measuring the value of financial compliance.

Goal	Metric
Reduce financial non-compliance risk.	Expected cost of failure, adjusted by the expected cost of remediation, plus the expected cost of preparation.

But wait – on the one hand I say "the value of compliance" and then the table says "non-compliance risk". Well, this is not inconsistent, because the value of compliance is that one has avoided non-compliance, and non-compliance has an *actual* cost when it occurs, and it has an *expected cost* before it occurs that is simply the product of the actual cost and the probability of non-compliance.[1] When planning for any kind of risk mitigation investment one must take both the actual and expected costs of failure into account. This has been explained in detail earlier in this chapter.

The point of Table 17 is that it gives IT a starting point from which it can start to build a model for how business losses can occur as a

[1] Financial compliance is not an all-or-nothing event. In practice, an accounting firm will generally find issues that need to be addressed before it will sign-off that the organization is in compliance. The seriousness of these issues factor into the magnitude of the "non-compliance".

result of IT system failures, assess their likelihood, and ultimately work toward computing a reasonable value for the risk of non-compliance. This provides a basis for planning appropriate-cost IT solutions that can mitigate that risk.

Once the metrics have been stated, it is important for the business and IT to work together to develop a model for computing each of these metrics. For example, Table 18 provides a possible formula for computing the metric for the value of additional products.

Table 18: A model for computing metric 1.

Goal	Metric
Introduce additional competitive products.	Expected value of product over ten years, adjusted for direct and indirect costs.
Model: NPV(Expected revenue MINUS (Expected direct costs.) MINUS (Integration costs.) MINUS (Maintenance direct costs.) MINUS (Marginal cost of adding additional systems.) MINUS (Cost of maintaining knowledge and skills.))	

Managers in operations – including IT operations – are accustomed to computing costs, and product and marketing managers are accustomed to computing revenue and the value of opportunities. The challenge is to put these together, and that is what Table 18 attempts to do.

If you have examined Table 18, then you will probably have noticed a few items as unconventional:

- Marginal cost of adding additional systems.
- Cost of maintaining knowledge and skills.

The first of these, the cost of adding additional systems, refers to the value of agility, as explained earlier in this chapter in the section Comparing Alternative.

Explicitly including the cost of maintaining knowledge and skills puts teeth into the truth that is widely known within IT, that as systems age they become harder to maintain because people with knowledge of those systems leave, and one has a choice of either taking steps to maintain that knowledge or bearing the cost of lost flexibility and ultimately a difficult and costly reverse-engineering or replacement process. By including these effects into the model, one is forced to make these hidden but very real costs explicit.

Modeling financial non-compliance risk is similar: one merely needs to identify the direct and indirect costs. This is not rocket science: it is politics.

Investment Preferences

Organizations – or any investor for that matter – do not always pursue the course that has the highest expected return. This is because investors usually have preferences for one kind of risk over another, driven by goals besides financial return. For example, many "green" companies seek to maximize profit, but subject to the constraint that the company does well by the environment or the human condition in general.

Even organizations that operate strictly for profit are biased in their preferences for one kind of risk over another. For example, some organizations are very conservative with regard to risk, while others are not. The oil drilling industry is accustomed to extremely large risk investment, and so it has evolved to deal with that risk. Organizations also tend to pursue investments that they are familiar with or that access familiar markets in order to reduce the risks and costs associated with acquiring new skills or new markets. Finally, some organizations have a preference for short-term investments while others prefer long-term investments.

Catastrophic Loss

One must always consider the possibility of catastrophic loss, such as loss of a business license or charter, an irreversible loss of market share, or loss of facilities or access to capital. Even if the expected gain is favorable, the probability of catastrophic loss must be considered because it would mean a "game over" situation. Many business plans merely consider the failure modes and refer to these as "risks", but a better risk analysis assigns a value to risk – even if the value is a magnitude or uses a relative but agreed-upon scale such as "low", "medium", and "high".

One way that consideration of catastrophic loss can be incorporated into the approaches outlined above is that one can determine the level of mitigation that reduces the probability of all kinds of catastrophic failure to an acceptable point, and set that level of mitigation as a minimum: the optimal, as determined by the analysis explained above, must produce an investment level that is greater than the minimum; otherwise, the minimum is used.

This is a simplistic approach, and does not account for the fact that even conservative organizations (or people for that matter) will sometimes accept some possibility of catastrophic loss if the up-side is high enough. In fact, everyone does this every day, whether they want to or not, since the ever-present possibility of catastrophic loss is part of life.

Preferences Given Limited Investment Resources

Another issue is that total investment dollars are usually limited, and so one must consider how investment in mitigation might reduce funds available for investment in other opportunities. Thus, instead of maximizing the return from a particular kind of mitigation, one must attempt to maximize overall return to the organization, constrained by limited funds, taking account of the organization's preferences for risk versus opportunity and for the makeup of its investment and project portfolio, as well as preferences for one kind of risk over another.

It turns out that when one takes account of organizational preferences such as aversion to risk, then the global optimal for the organization will be the point at which the "marginal rate of substitution" (of one kind of investment over another) equals the "marginal rate of transformation" (of one investment into another). Readers who are interested in these concepts in depth are referred to [Copeland03] and [McAfeePr06]. Here I will provide a brief discussion of how these concepts affect our models.

Indifference Curves

Consider an organization that provides two services, A and B. These services may be any kind of service that decision-makers within the organization feel are desirable for the organization's constituents. For example, service A might be the production of a particular product for sale, and service B might be the production of another product for sale. However, in the discussion here it will not matter what services A and B are, and so I will simply refer to them as A and B.

The organization will, through its decision-making processes, exhibit a preference with respect to the relative value of A and B. This preference might result purely from an attempt to maximize shareholder value (for a for-profit corporation), or it might reflect other preferences, such as a desire to focus the organization's activity more in one direction than another, or both. Assume that the organization is currently producing a certain amount of A and a certain amount of B. If one were to propose decreasing the amount of service A being produced, by how much would one have to increase the amount of service B being produced in order for executives to say that the overall business value remains unchanged?

The answer to this question defines a point on a curve, known as an *indifference curve*, because the organization is indifferent to where it sits on a given curve: that is, if it is at one point on such a curve, then it is happy to move to any other point on the same curve, by definition. The position of the curve defines a level of "business value", or "preference".

If one fixes the budget for the organization, so that a fixed amount of funds are available for both A and B, then the allocation to A and B is defined by a curve also. This curve is known as the *transformation curve*, because it defines, for a given budget, how much of A and B can be produced. This curve is usually convex, but for simplicity I will treat it as a straight line: that is, I will assume that if one spends twice as much on service A, then one actually gets twice as much output from service A; and similarly for B. Also, I will refer to this curve (straight line) as a "budget line", which I feel is more intuitive. Consider Figure 20.

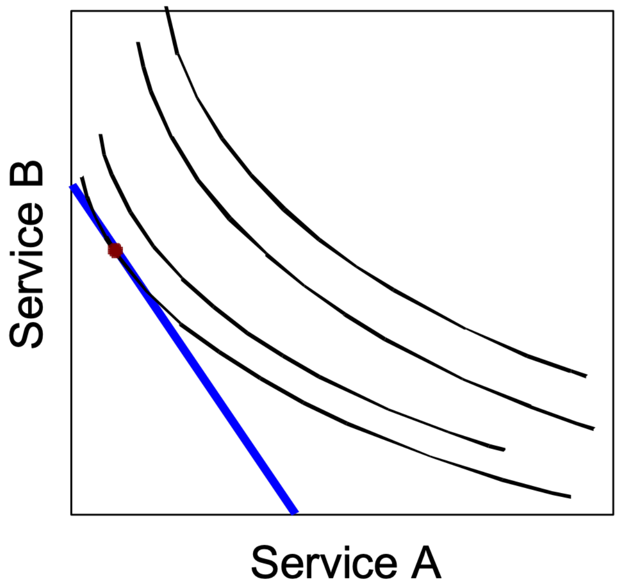

Figure 20: Indifference curves and a budget line.

Figure 20 shows a straight budget line and several indifference curves. Each curve represents a different level of perceived business value: the higher the curve, the higher the business value. The budget line shows how much of A and B can be produced for the fixed budget available to the organization. If the available budget were higher, the budget curve would be higher.

The point at which the budget line is tangent to an indifference curve is the point at which the organization reaches the highest possible indifference curve, for the given budget. Thus, this is the point at which the organization should divide funds between A and B.

Remember, an indifference curve is a curve along which the perceived business value is the same at all points, and represents tradeoffs between two services. (For three services, one has an indifference surface; for four, an indifference hyper-surface, and so on.)

A reminder of how this relates to maximization of business value: if we assume a given budget for investment, and optimize our investment strategies and portfolio, we should end up at the point of tangency indicated in Figure 20. The purpose of Figure 20 is to show graphically how the budget affects investment.

Indifference curves are of most utility when one has to balance non-financial objectives, because if the objective is simply to make the most money in a given period of time, then one need not consider preferences at all: one simply chooses the strategy that yields the greatest expected return. On the other hand, risk aversion (safety) is a kind of preference, and so indifference curves can be used to model preferences for safety versus possible gain.

When Business Value Is Not Financial

The mission of a non-profit organization or a government agency is not to make money. Therefore, for such an organization, one cannot assume that one can measure business value in terms of money saved or money earned. However, if one assumes that the organization's budget is optimal, based on the willingness of its constituents to invest, then one can assume that *the marginal value of each additional dollar spent equals the economic value (to the constituents) of that dollar*. Given this thinking, a dollar saved releases a dollar for investment in a dollar's worth of additional services, according to the mix of services that are provided by the organization. Optimality assumes that this mix represents the preferences of those who fund the organization's services: its constituents, or the taxpayers in the case of a government agency.

Mandates

A mandate from a regulatory agency, from law, policy, or any other source must be followed, and so it represents a constraint on strategy and decisions.

A mandate has no business value, because it is externally imposed, and no business value should be attributed to it. However, the actions compelled by the mandate might have business value, and the very strategy that is behind the mandate – the reason for the mandate – might have business value.

A business value analysis should therefore not place a value on an externally imposed mandate, but may assign a value to elements of a strategy or plan that implement the mandate.

Direct Cost Avoidance

A great many internal investments within organizations are made for the purpose of saving costs. For example, an investment in a new IT system might save the costs associated with the current paper process.

Direct cost avoidance such as this represents a net financial gain for the organization. As explained earlier, that gain can be re-distributed across the various services based on the organization's priorities.

Overall Mission Effectiveness

Organization governance, directives, and oversight processes are intended to ensure that projects reflect the global needs of the organization and not merely the narrow needs of particular services. If the organization has a parent organization (as all government agencies do), then overall effectiveness therefore has to do with the overall effectiveness of the organization part of the parent organization. Overall effectiveness includes areas such as fiscal responsibility, adhering to parent organization policies or regulations. Fiscal responsibility includes adhering to standards for IT, such as enterprise architecture guidelines, that are intended to promote cost effectiveness, reusability, accountability, security, and agility.

Overall mission effectiveness is optimized by finding the point at which a new budget line is tangent to an indifference curve for the organization.[1] This is illustrated in Figure 21.

[1] Those who are familiar with Pareto curves will note that I am assuming that the organization acts as a single decision-making entity.

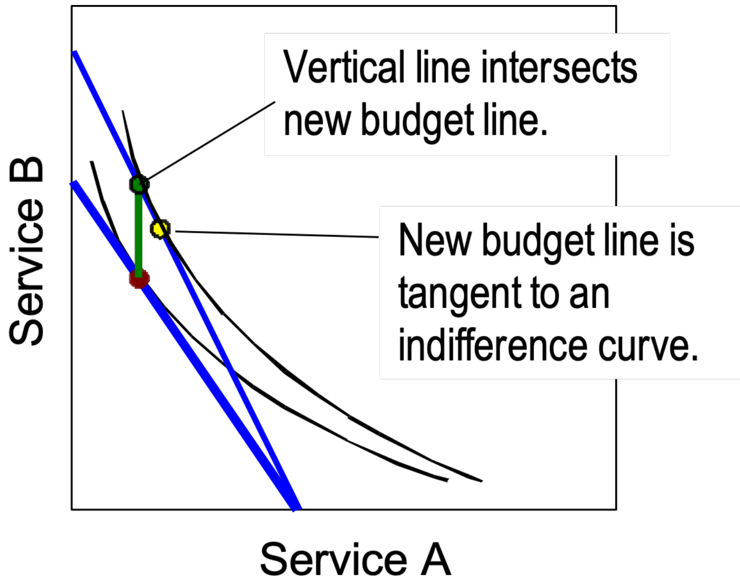

Figure 21: Command and control: When the preferences of only the organization as a whole must be accounted for.

A subtle point is that improving the efficiency of a service changes all of the indifference curves for the organization. This is not an important issue for the discussion here though. The main point is that if a service is made more efficient, management should re-evaluate the distribution of funds across all of its efforts to maximize the overall benefit to the organization. In most cases, this will result in maintaining the funding of the more efficient service, and possibly increasing it, since now more of it can be obtained for less, and so it is more cost effective.

Societal Benefit

Government agencies exist to benefit society. However, the quantitative benefit to society is different from the quantitative benefit based on the agency's mission. The agency has a limited budget and that budget is based on priorities defined by the government's policy makers and executives. Just because an agency service activity saved a million dollars for some segment of the general public (e.g., by protecting property during a natural disaster) does not mean that the activity was worth a million dollars in the context of the agency's budget. If might have cost $100,000 to save the million dollars, so when compared to the agency's budget items the activity was worth $100,000, if one assumes that the agency's activities are balanced across the agency's priorities.

It is therefore important to not impute value to public agency activities based on the direct impact on the public or certain constituents. Internal dollars do not equal external dollars.

A Real-Life Example of Estimating Tangible Value

Case In Point

The business proposal for a $5 million project made the following intangible claims of business value:

- Improved mission effectiveness.
- Improved tracking and reporting.
- Improved auditability.
- Improved transparency of business rules.
- Reduced oversight costs.
- Improved data quality.
- Improved security.

In addition, we were interested in the possibility of benefits related to reuse and lifecycle cost, so we asked questions about those areas.

The proposed system was an advance-procurement system, and its purpose was twofold: (1) to improve the availability and quality of equipment, so that right equipment would be available when needed, and the equipment would be of the required quality; and (2) to reduce the costs associated with equipment by allowing personnel to pre-negotiate equipment leases in advance of when it was actually needed, based on competitive solicitation.

I assess tangible business value for each of these in the following ways.

Mission effectiveness. Mission effectiveness had to do with the ability to perform using the equipment procured through the advance lease agreements created with the system. The ability to perform was impacted by the availability and quality of equipment. We modeled the value of mission effectiveness by estimating the *percent improvement* in effectiveness that was expected to result from having advance lease agreements for quality equipment. This was based on the actual improvements seen through use of a prototype of the system, and the actual impact on mission effectiveness. The estimate of improved effectiveness was obtained from a credible field expert. By estimating improvement as a percent, we were later able to normalize the value of the improvement to the cost of current operations: if current operations in the area of business cost X, and effectiveness was improved by 10%, then 10% of X is free for re-distribution across the organization. Even if most (or all) of that 10% is kept by the area of business, the value to the organization is the value of having .1X more funds available for investment.

Improvements to tracking and reporting. The current system – to be replaced by the new system – was paper-based. This resulted in delays in payments on lease agreements, resulting in interest charges that had to be paid out. Thus, a decrease in turnaround of payments translated into reduced interest payments to vendors. We asked experts to estimate the decreased frequency and amounts of late payments as a result of automating the paper process, and computed as estimated reduction in interest charges. Since the current paper system processed hundreds of million of dollars of transactions per year, and payments were often late, the interest savings were substantial.

Improved auditability. For this system, financial auditability was required by law, and so we treated it is a mandate.

Improved transparency of business rules. The prototype system had little transparency with respect to the business rules encoded in it, and the business users therefore had concern over whether the rules were correct. Therefore, for the new system, a promise was made to the business users that the rules encoded in the system would be transparent and could be reviewed. The value of this was therefore treated as a mandate.

Reduced oversight costs. The current paper process was spot-audited at great cost, requiring analysts to travel on-site to review paperwork. The new system would make it possible to review data electronically from the oversight auditor's office. We estimated a savings in travel costs and time based on projections from an expert (in the oversight group) about the reduced travel and reduced time to audit data.

Improved data quality. A prior report had documented a 50% correction rate for the paper lease agreements that were in use. We extrapolated improvements obtained from the prototype system to estimate the reduced corrections required for the new system, and the cost savings that would result.

Improved security. The prototype (which handled 10% of the volume that the final system would handle) had experienced six documented security incidents in three years. One of these resulted in a severe loss of reputation because the head of a vendor association wrote a memo in the association's newsletter urging vendors not to use the prototype system. This action threatened the entire value proposition of the prototype system. By extrapolation, if the new system were no more secure, then this kind of event could be expected to happen even more frequently given the much higher volume of processing of the new system.

Reuse. The project team insisted that they expected future reuse of their system's components to be zero, even though it was built on a highly reusable SOA architecture, based on their experience with cross-departmental competition and trust. This issue was in fact an issue that was being worked at a higher level, and so it was not a surprise, even though it was a disappointment. We assumed zero benefit from future reuse.

Lifecycle cost. The project insisted that given their rural geographic location, that they could expect a high level of continuity of the team over time, and that when the system eventually went into a maintenance mode, they would be able to transition some of the developers to maintaining the system. In this way, they expected that the valuable knowledge acquired as a result of building the system would be leveraged and not discarded. We estimated that 20% of the team would remain to maintain the system, and projected a substantial reduction in maintenance cost and improved reliability and extensibility as a result.

Modeling the Interactions

All of these claims for value could simply be added, to produce a total claim for value. However, there is a more powerful way to estimate overall value: by modeling the sources of value and their interactions. That approach enables one to perform sensitivity analysis, to determine where it would pay to invest more, or which sources of value depend on sufficient investment in those areas. For the project under discussion, I developed a model, as shown in Figure 22.

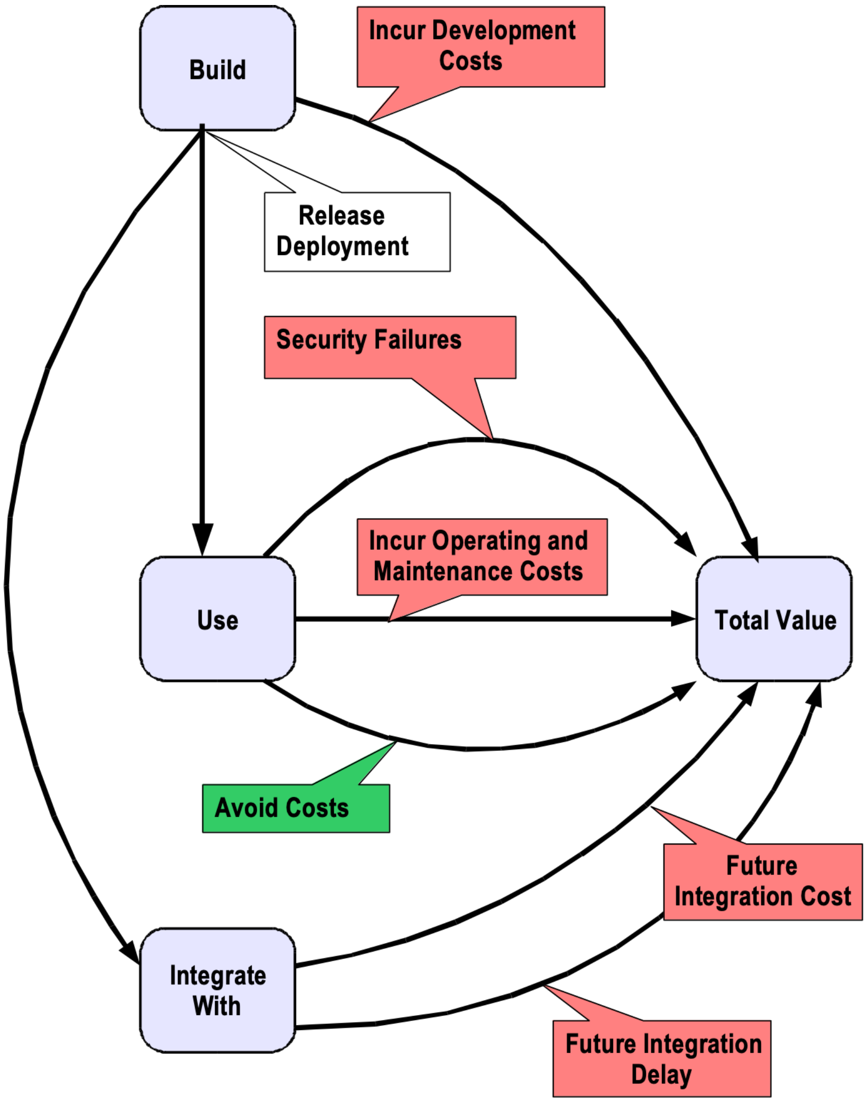

Figure 22: Cause-and-Effect Value Model for the New System.

This model contains representations of the various sources and detractors of value, and their interrelationships. For example, building the system incurs a development cost, and this is shown by the arrow leading from "Build" to "Total Value". The arrow carries a negative value whenever a development cost is incurred. Usage of the system, once deployed, is represented by the "Use" bubble, and the costs of use – operations and maintenance (O&M) costs – are represented by the straight arrow leading from "Use" to "Total Value". Again, this arrow carries a negative value whenever O&M costs are incurred. On the other hand, the lower curved arrow leading from "Use" to "Total Value" represents the business value produced by the system's function – its reason for being. The values conducted along this arrow are positive.

The bubble labeled "Integrate With" is perhaps the most interesting and subtle. It represents hypothetical future integrations with the system. Whether these future integrations will occur depends on the nature of its function, and also how easy it is to integrate with. Assuming that future integrations will occur, their cost will depend to some extent on the way in which the system is built today: that is, whether the design of the system makes integration easy or not. The upper arrow leading from "Integrate With" to "Total Value" represents the costs of future integration events as they occur: this cost will be less if the "Integrate With" bubble assumes a strategy of using an architecture that is easy to integrate with.

The lower arrow between "Integrate With" and "Total Value" represents the time-to-market cost of the time it takes to implement and deploy future integrations sooner. For example, if a future system requires integration with this system, and the architecture of this system is such that the integration can be done six months faster than if a traditional architecture were used, then the time to integrate will be six months less, and so the new system will be deployed six months sooner. Presumably there is business value in having the new system deployed sooner, and the lower arc represents this value.

In addition, the model accounts for security-related failures, that is, incidents that cost the organization in some way. Security-related failures are modeled by the arc from "Use" to "Total Value" labeled "Security Failures". The costs that travel on this arc represent the cost of each security-related incident as it occurs. These are negative sources of value. Also, these follow a distribution that has a very broad range, from minor incidents to incidents that undermine trust in the system by other organizations and partners and therefore threaten the entire system.

The model was simulated, using a variety of assumptions for the time distributions and cost distributions represented by the arcs in the model. This made it possible to see how sensitive the model was to the various assumptions, and it gave a range for the total business value to expect over time.

It turned out that the total value was very sensitive to the degree of future integrations expected and their value. The implication was that we needed to pay particular attention to projecting whether such future integrations would actually occur, and what their value might be.

The model results are also sensitive to assumptions about the magnitude of catastrophic security-related incidents. The implication was that investment in security risk mitigation to present catastrophic class incidents is highly recommended. In this way we were able to estimate the "value of security".

Summary

In order for IT architecture choices to be taken seriously by business decisionmakers, those choices must be expressed in tangible business terms.

Business value results from expected income from planned sources as well as expected income from future – and as yet unplanned – sources. Both of these must be considered in some manner, because both are impacted by architectural decisions today. Busines value is also (adversely) affected by expected investment costs and expected lifecycle costs.

Architectural decisions that provide future flexibility produce business value by increasing the expected income (or value) from as-yet unplanned sources. Thus, the value of agility is directly related to flexibility and to the expected magnitude of the unforeseen opportunities that can leverage that agility.

In order to prove the value of business agility, one must measure the enterprise-wide impact of decisions made years back. Therefore, it is important to retain business effectiveness metrics over a long period.

Risk manifests as an expected loss due to foreseen or unforeseen types of events. Investments in risk mitigation increase value because they reduce the size of the expected long-term loss.

There is an optimal point at which the marginal value of investment in risk mitigation is zero, and that is the point of maximum value with respect to risk mitigation.

Decisions about risk mitigation and investment in top-line improving activities should be compared together on a common scale.

In IT it has been especially difficult to measure the value of system features because many features do not operate independently and so they produce no direct business value of their own. Rather, their value derives from their support of other features. This is problematic because it is then difficult to estimate the value of completing the development of these features since they are so inter-dependent. This can be resolved by focusing on the increased *expected value* of the total set of features. This expected value must account for the risk of completion of the full feature set and how that risk and cost is reduced as a result of completing embedded features. The value must also accunt for the opportunity that is presented once all featuers are complete.

Very often architectural alternatives must be compared in which one alternative provides more flexibility than the other. Estimating the value of this flexibility is difficult. It can be done however by hypothesizing a future change that takes advantage of the flexibility and estimating the opportunity value of that flexibility as well as its likelihood. In this way an *expected value* of the flexibility can be estimated.

Progress in building software is actually a statistical phenomenon: it is (a) the increased chance of finishing and thereby realizing the future value of the new capability, plus (b) the reduced expected cost of completion. Project management is therefore all about probabilities, judgment, and risk management.

It is not practical to perform detailed value analysis for every IT decision. However, it is important that many kinds of decisions – such as those that affectt security or other kinds of risk – be made in a consistent manner that reflects the relative risks. Toward that end it is useful to develop a decision framework that enumerates practices and provides risk-based criteria for when to use each practice. In this way the risk and value analysis is done when the framework is created rather than each time a decision is required.

Knowledge acquired as a result of a business activity is a valuable asset, and this value must be tracked and taken into accunt when decisions are made regarding how to deploy resources. Knowledge should be tracked as an asset and should be part of the organization's meta data. The value of knowledge depends on the task, and that value is the amount by which it increases the expected value of an investment in pursuing the opportunity associated with the task.