Pittman
CONSULTING

A Landscape of Risk

Project Risk Management

This version published electronically by Project Managers Network with the permission Pittman Consulting.

July 1999

Author:	Matthew Pittman
Date:	29 June 1999
Version:	1.00
Status:	Final
Document ID:	PCPL-ART-002

Introduction

I had almost given up hope. Was it just a flaw in the way I expressed the concepts or was I cursed with the project management equivalent of the Cassandra complex, destined to predict the future but not be believed? OK, perhaps it was not as melodramatic as all that but I had never had much success convincing clients that risk and its control was key to successful project management, until now.

It turns out that the old adage is true, "You can't tell someone something they don't already know".

Now, the confluence of three trends has informed my audience for me and I no longer need struggle to explain risk in projects - clients are explaining it to me!

The first trend is the increased awareness of risk in financial markets brought about by the Asian crisis, the Russian debt default and their precursor, the collapse of Barings Bank. This sharply focused the interest of the entire business world on the damage that can be wrought by unmanaged risk. No convincing argument works so well as a cataclysm of mythical proportions, starring a hero with a tragic flaw.

The second trend is the wonderful focusing effect on management of the drive to deliver shareholder value when combined with a swag full of share options. Management's attention to the calculation of value has brought the lesson of Nick Leeson right into the boardroom. Bless you, Nick, for you have added to every comparative process the need to add a dash of "regret" analysis here, a sensitivity analysis there and a scenario or two.

Finally, the millennium bug has made software risk management, its costs and its benefits a central issue. Although I feel the industry may live to regret this effect.

So now when I talk with clients about project risk they are all ears.

The cover of a book on risk I read recently captivated me. Seeing Tomorrow by Dembo and Freeman shows a person overlooking a valley shrouded in mist.[2] I will borrow that image as the theme of this short article, a landscape of risk. I want to discuss risk as it relates to project management. I want to illustrate how we can view risk, not just as a feature in a project panorama but as the defining force in the formation of a changing landscape - a complex contour of risk plains and peaks that must be understood to be safely negotiated.

In my mind's eye, no feature of this landscape is fully formed. Like some shimmering quantum world, a range of probabilities defines all the features of the landscape. No feature becomes clearly defined until the risk surrounding it is controlled. Indeed the features remain hazy until the passage of time finally decides what form they should take. They flash into focus for the briefest instant and then quickly fade into the mists of memory, achieving clear definition only at the moment they are encountered.

I have also set myself two additional goals, to deliver a few, simple, practical techniques for helping improve our view of this landscape and to illustrate the potential flow of risk management ideas from project management to business and finance and back again.

There is insufficient space to deal with all aspects of project risk or risk management; that would require a book. My intention is to deal with risk and uncertainty in different corners of our project landscape, estimate, schedule, bid, project and financial risk.

Differentiating Risk and Uncertainty

Uncertainty and risk are two different but interrelated concepts that are sometimes confused. Many aspects of a project may be uncertain but they may not be risky.

Risk is the possibility of danger, loss, injury or other adverse consequence. Uncertainty is not knowing for sure.[1] Being uncertain about a condition such that the condition may possibly result in loss is a risk. Just how risky it is or how exposed to the risk we are depends on the answers to several questions.

- How probable is the outcome that leads to loss?
- How much loss might be incurred?
- How much will sufficient[2] knowledge about the condition cost?
- Is there anything we can do to avoid or reduce the impact of the risk and, if so, how much will that cost?
- How important is it relative to other risks?

There can be no risk without some uncertainty since, if the outcome is certain, the loss is no longer a possibility but a fact. Some risks are so unlikely, so catastrophic or so outside our control that they are effectively of no interest. Thus in projects there can be a great deal of uncertainty and some risk. Our purpose here is to deal with the risk.

The process of identifying those risks that are meaningful in the context of our project and answering the questions, above, about them has been referred to as Risk Assessment. The process of doing something about the risks has been referred to as Risk Handling and together these constitute Risk Management.[4]

Estimate Risk

Estimates are the perfect place to draw a line between risk and uncertainty. Whatever technique is used to arrive at estimates it will have an element of uncertainty, which may vary for each individual estimate. These uncertainties may present a risk to the project or they may not.

If a project manager thinks the estimates do represent a risk then the well-known method of applying contingency is still the most appropriate. Contingency is a way of mitigating estimate risk and can be calculated per item or for the estimate as a whole. Since uncertainty in estimates impacts the schedule, some options for allocating contingency will be discussed in the Schedule Risk section, on page 4.

My software project manager colleagues and I seem to agree that, in the absence of any better analysis of the estimate uncertainty and other schedule-related risks, the normal contingency to apply should be around 30% of the unadjusted estimate. A project of 1000 person days would have an estimate contingency of about 333 person days. I have never seen much empirical evidence to support this position - it is a rule of thumb and it assumes that the project manager will be dealing with other risks explicitly rather than relying solely on contingency in estimates.

No matter how we may wish that estimates for projects could be made before designing a solution to the problem the project is trying to solve, the truth is that the vast majority of estimates are based on a feasible design (if only by analogy). The software industry seems to be alone in thinking that it should be possible to estimate the cost of something without knowing or guessing at the something's design characteristics.

I suppose there are limited circumstances where a raw statement of requirements can be used for estimating, such as where substantial data is available on previous efforts and it is reasonable to assume that the new project conforms to such

[1] after the Concise Oxford Dictionary

[2] enough knowledge to reduce uncertainty to an acceptable level (a subjective concept)

previous efforts. There are estimating techniques that support this style of estimate (Function Point Analysis and COCOMO II) but even these techniques produce more certain results if the requirements are developed into a feasible design.

When we develop a feasible design for the purpose of estimating we are likely to introduce another form of estimating risk. In making a particular set of assumptions and decisions about the feasible design we introduce the potential that these may be wrong. Correcting such a mistake can be extremely costly and so risk management techniques must be used to control the risks.

An effective technique that I use for mitigating the risk of getting assumptions wrong is to make sure that the assumptions and dependencies for the estimate are explicit and documented. Sounds obvious, does it not? In fact it can be quite hard to find a balance between documenting too many assumptions (such as that the sun will rise tomorrow) and too few. It is somewhat in the nature of assumptions that we tend to take a great many of them for granted.

The impact on the estimates of violating the assumptions or not meeting the dependencies is also documented. For example, if we assume something about the platform for the project we may want to consider the impact of having to select a different platform. Assumptions and dependencies are also discussed in Bid Risk, on page 5.

Finally, it is always important to estimate the cost of risk management. It may cost quite a lot to answer the question about the platform. In fact it may cost more than it would to change the platform. There is more on this issue in Project Risk, on page 6.

Schedule Risk

The schedule combines task estimates, resources and dependencies into a plan that delivers the all important project end date and project cost. This is the project's bottom line.

Uncertainty in the estimates introduces risk to the schedule. A single uncertain estimate is unlikely to have much impact but if all the estimates for all the dependent tasks in the schedule are uncertain, the effect is cumulative. This is why the introduction of contingency is so important. Contingency is designed to cover the cumulative uncertainty in the schedule.

A straight percentage of the effort or the elapsed time is perfectly adequate to manage most schedule risk. There are various effective options for handling contingency in a schedule.

- Distribute effort-based contingency through each task, prior to resource allocation.
- Distribute elapsed time contingency through each task, after resource allocation.
- Add elapsed time contingency at critical delivery points throughout the schedule or at the end of the schedule.

The reason for making the distinction between effort and elapsed time is simply that contingency for effort estimates is consumed in proportion to the resources allocated to it. We need to make sure we are adding effort contingency to tasks in proportion to the allocated resources and elapsed time contingency to the schedule as a whole.

Some project managers have elevated the management of contingency to an art form. They recognise and manage additional factors.

- When is the contingency "burned" (either by a successful delivery or by the extra work required)?

- How fast contingency burns? More contingency is burned early in the project because of relatively greater uncertainty compared with late in the project.
- When and how to show contingency to the team and senior management?

The schedule introduces another risk. Tasks and dependencies combine with the level of uncertainty in the estimates to produce unusual and unexpected effects on the critical dates. The project delivery date can be quite volatile as a result of this complex interaction and in complex schedules it can be difficult to predict these effects using static analysis basically because relatively small changes in estimates can have disproportionate effects on the project's critical path.

One technique that is used to assess this schedule risk is to simulate changes in the estimates according to a specific probability distribution and analyse how these impact the critical dates in the project. This technique is known as Monte Carlo Simulation.[5] In effect it provides a large number of scenarios for sensitivity analysis.

A Monte Carlo simulation can be a great tool because it enables the project manager to provide a risk adjusted project duration and cost. However, the only risk covered is the schedule risk and estimate risk, not project or other risks.

For a Monte Carlo simulation to work you have to know or guess the probability distribution for the estimates (in some cases there may be different distributions for different types of estimate). This means you must either keep meaningful data from past projects (something most organisations still do not do) or manage the risk of guessing wrong.[3]

Another common problem with Monte Carlo Simulation is that the schedule has to be a good one. By which I mean it has to be detailed and include correct dependencies. Obviously running a simulation on an inadequate plan will yield an equally poor result. I admit to musing about the possibility of using Monte Carlo Simulation on any old schedule and using a volatility measure of the result as a measure of the "goodness" of the schedule.

One common misconception is that Monte Carlo Simulation will give you a single answer to the question of project cost and duration. I suppose that if you take the worst case result from the simulation then you could say that but the result is likely to be a plan for an extremely costly and very long project. What you do get from a Monte Carlo Simulation is a distribution of results that will allow you to predict with known confidence an outcome in a certain range.

A simple technique that does not require as much effort as the Monte Carlo approach is to record best, nominal and worst case for each estimate. You can use a spreadsheet or the extra fields in MS Project to record these numbers and adjust the schedule as a whole manually or using simple macros.

Bid Risk

Bidding jobs is part of winning business and often the bids require estimates and plans in short order. There are two competing forces in bid work that act like tectonic plates forcing past one another. The movement in this landscape is along fault lines of risk, which separate sales and project management.

The risks that all project managers immediately identify in bid work are the estimate and schedule risks. These are often compounded by the time pressure of the bid.

---

[3] There are plenty of tools for Monte Carlo simulation that make understanding the probability distribution a simple task if you are prepared to make a few guesses about the uncertainty in your estimates.

The risk that the sales and sales support people identify is the risk of losing the business because risk management makes the bid uncompetitive.

Worst case risk management will not work effectively here because of the risk of losing the business. For the project manager to achieve a safe bid requires a combination of techniques to help mitigate and offset the risks while still maintaining a feasible commercial position. Here are some tricks of the trade that may help.

- Make contractual assumptions and dependencies that expose major risk items and exclude them from the contract.
- Analyse the risk in the project using statistical risk analysis based on simple techniques such as the Risk Parade described in Project Risk, below
- Vigorously define and control project scope using explicit inclusions and exclusions.
- Trade high-risk items for lower risk items during the bid negotiation phase by taking risk into account for the final price of each line item.

Project Risk

In the day to day running of the project much of the project manager's time should be taken up with risk management. Most project managers are engaged in a fair bit of "fire fighting" on a project. Many project managers do not see this as risk management but it is. One trick is to try to make this fire fighting a more explicit risk management process.

The best project managers seem to be the ones who can get to fires and put them out before they get out of control. They seem to know intuitively where the fire hazards are. In other words they know about risk.

Making risk management a normal part of the project manager's activity helps to refine his or her sense of risk and gradually they get much better at what used to be ad-hoc fire fighting.

There are lots of tools around to help with risk management but the best one I can recommend is so simple many cannot believe it works but I assure you it does.

Someone came up with the idea of keeping a list of the Top Ten Risks on a project. I think it was Barry Boehm but it could have been Tom Gilb or someone else, I cannot remember. When I started using the technique I found it kept my mind focused on the important issues in risk management. Gradually I took on board other ideas I picked up and refined the technique until my Top Ten Risks became my risk register and main risk management tool. I named it the Risk Parade in honour of the original idea (whosever it was).

In this instance it can be best understood as a series of refinements, each one has a reason for being there. Use this as a cookbook to allow you to make a habit of risk management. Do not be tempted to put more features into it at the start of a project that you will not be prepared to maintain at the end.

1. I used a spreadsheet to make a list of all the risks that the team and I could identify on the project. I started with a subjective ordering of most severe risk to the least severe and I kept a record of the position at the last risk report or status report. This was the "Top Ten" risk list.
2. The impact of a risk is impossible to assess if we do not know its likelihood and its impact. I made a subjective assessment of the probability of the risk occurring and I estimated the impact. I found that the easiest way to estimate the impact was to estimate the cost to fix the problem if the risk "fired". This quantity was named "cost to fix".

3. In order to rank the items in the risk parade I multiplied the probability of the risk by the cost to fix to yield an exposure for the risk. At this point we have all that is needed to have a functional risk parade. Project Managers can use this list to record risk and prioritise them for action. More useful still is that the total of all the exposures is a nominal value for project risk exposure for the project. However, this exposure is meaningless if there is simple action we could take to reduce the exposure.
4. The next trick is to add a column to include a description of the mitigation that is planned. This helps make the planning of mitigation and its execution explicit.
5. Of course to assess the effect of mitigation we must estimate its cost and its impact. In this way it became possible to assess the impact of mitigation on the risk and compare it against its cost. The risk that remained after mitigation was named residual risk. If mitigation plus the residual risk cost less than the unmitigated risk then the mitigation was a good candidate for action. The total mitigated exposure combining the residual risk and the cost of mitigation represented a genuine measure of the likely exposure of the project to risk.
6. There were many more enhancements that I tried at different times. I found it valuable to track the costs and exposures in terms of person days, elapsed days and other costs ("on" costs, etc). Person days and other costs were used in the calculation of financial exposure, which was used to rank the risks. The impact to the schedule in elapsed time and effort was a separate exposure. A colleague added a start and expiry date for the risk so that it was easier to determine when a risk was no longer in effect.

This simple system is still more advanced than many of the risk management tools I have tried to use but the Risk Parade only works this way if the risks can be made independent events. If this is not possible then the exposure will be miscalculated. The way I "cheat" to get around this is to calculate the dependent probabilities and record the risks as independent events with adjusted probabilities. For example, if there is one chance in ten of a problem with a new product and one chance in ten of that problem being "terminal" then the risk entry would have two independent events, a 0.10 probability of a product problem and a 0.001 probability of a terminal product problem.

Financial Risk

One of the themes of this article is how the view of risk that is now accommodated in general business flows to risk management in projects. Risk management in projects then influences how company management views the project so the influence flows both ways. This is particularly so now that "management by project" is becoming more common.

Managers are expected to make decisions based on sound financial analysis and projects compete for funds against other projects. Generally this means that a project manager must provide a persuasive financial argument for a project and this argument must enable the project to be compared with other projects in an objective manner. Project managers have learned the techniques for presenting the finances in these business cases from the business techniques used in strategic evaluation.

Today, the most accepted strategic evaluation technique is Discounted Cash Flow (DCF), comprising the production of figures for Net Present Value (NPV) and/or an Internal Rate of Return (IRR).[1,6]

NPV works out the value of a project and takes into account the fact that a dollar returned after a project completes is worth less than a dollar invested in the project today. This makes sense because you have to obtain the investment dollar from somewhere and getting this capital costs money (interest on a loan or dividends on

equity). The money that we wait for in return cannot be used until we get it and this also costs money.

When we want to compare projects we can compare their rate of return (their profitability). For any investment there is a value for IRR that represents the rate of return where the returns balance against the costs. For a project to be profitable it must have an IRR greater than this "hurdle rate". Obviously IRR is handy for comparing two projects because NPV alone gives no indication of the return on our investment.

Many project business cases use these techniques now and project managers have readily adopted them because they are generally simple, transparent and objective. NPV and IRR are supported by spreadsheet functions so they are not hard to calculate. However, NPV and IRR alone do not take account of risk. Given that it is very likely to be a project manager working on a project's business case this is a real problem.

Comparing IRR without adjusting for risk is like using NPV to rank different projects. The financial analysts realised that NPV gives no indication of the rate of return on an investment. In the same way IRR does not make visible the risk we are taking with an investment.

Most business cases that do take account of risk do so using a well-known technique. The major percentage rates and other inputs are varied within an expected range and the results are analysed to see how NPV and IRR vary. This yields a best, worst and expected value for NPV and IRR. Most spreadsheets are capable of a limited form of automation in this area.

The same technique is often used to see how cash flow changes when the inputs change. This is sometimes referred to as Sensitivity Analysis.

These techniques help mitigate the risk of incorrect assumptions, changing market conditions, project delays and the like.

Different decision rules that can be applied to these results help introduce risk management to the question of picking projects in which to invest. The obvious ones are optimistic (best case) and pessimistic (worst case). More interesting is the idea of measuring the cost or the "regret" of making the wrong choice and then attempting to minimise the regret.

Regret has been introduced to decision analysis for financial market investments.[2] Regret itself is not a new idea but using it as part of a risk-adjusted return on an investment is a recent application and I think there may be some lessons there for business case presentation.

This idea of a risk-adjusted financial analysis for projects should be attractive to project managers. After all, we strive to quantify risk so that the effects can be seen in the final figures, so it makes sense that we try to develop a technique for expressing the risk in a way that can be used as input to financial decision making. In other words you cannot look at a set of NPV and IRR figures for three projects and make a decision because you do not have visibility of the riskiness so you cannot judge how volatile the outcome might be. You need some form of risk-adjusted return for the projects.

Sensitivity Analysis goes some way to providing a measure of volatility but it does not actually provide risk adjusted values. The difficulty is in providing a benchmark against which volatility or risk can be measured. This is easy in financial markets because there are benchmarks everywhere (cash rate, S&P 500, Dow Jones, etc). However, it is not so easy in a single project because there is no benchmark and there is not a set of return values that can be used to calculate a risk-adjusted return.

As a starting point we could compare the nominal project figures with the best and worst results from the sensitivity analysis. In a similar way we could compare the

unadjusted project figures, those where no account was taken for any risk (estimate, schedule, bid, project) with the risk adjusted figures. Either of these techniques should give us an indication of the volatility of the project outcome and from there it's easy to see what the risk adjusted return should be.

Monte Carlo Simulation of the risk factors affecting the project (including financial sensitivity analysis) can be used to provide a genuine measure of risk adjusted return for a project. There is not enough space to go into that here but I will add that a prerequisite for such an approach is an explicit and open handling of risk right across the project landscape.

Conclusion

Believe it or not, once you start with basic risk management the sophisticated analysis of the Financial Risk section, on page 7, seems neither daunting nor a waste of time.

It all begins with a basic understanding of risk and uncertainty in projects. From there it is easier to understand the risk landscape in estimates and in the schedule. It explains the competing forces between teams in the bid process.

There are both simple and advanced techniques to help manage the risks in each corner of our risk landscape. Where there are no techniques apparent they can be developed by the same evolutionary process outlined in the steps, on page 6, and in the final part of the Financial Risk section.

Whatever else a project manager may think of risk management, his or her perception of the project stretched out before him or her will change forever when simple tools are used to expose and explore a landscape of risk.

References and Further Reading

1. *Exploring Corporate Strategy*, Text and Cases, Gerry Johnson & Kevin Scholes, 3rd Edition, Prentice Hall, UK, 1993.
2. *Seeing Tomorrow*, Rewriting the Rules of Risk, Ron S. Dembo & Andrew Freeman, John Wiley & Sons, New York, 1998.
3. *Software Engineering Economics*, Barry W. Boehm, Prentice Hall, New Jersey, 1981.
4. *Software Risk Management*, Barry W. Boehm, IEEE Computer Society Press, Washington DC, 1989.
5. *Statistical Methods for Software Quality*, Using Metrics for Process Improvement, Adrian Burr & Mal Owen, Thomson Computer Press, UK, 1996.
6. *The Portable MBA in Finance and Accounting*, John Leslie Livingstone, John Wiley & Sons, New York, 1992.

