A Project Risk Metric

Robert W. Ferguson
Software Engineering Institute

A risk metric is proposed that is normalized across projects. The purpose of the metric is to provide management visibility into project uncertainty. This works best in an organization that manages multiple projects. The proposed metric can be applied early and throughout the project. It has been useful for identifying or canceling projects in trouble. It has also been useful for identifying projects that do not yet have a satisfactory risk plan.

The Standish Group published its original "Chaos Report" [1] in 1994 declaring that American companies spent $81 billion on cancelled projects. Additional Standish Group data in Figure 1 shows that the situation has not improved as much as one would hope.

Even projects that are not cancelled may deliver such reduced functionality that most people would not count them as successful projects. Often there has been early evidence that the project was headed for disaster. The project manager may even have issued warnings to senior management or sponsors about the problems. There simply seemed to be no way to *pull the plug* until the project was already over budget, late, and at the point where the customer was ready to give up or worse.

The problem may be a failure to examine the risks of the project from a systemic view. When risks are faced one problem at a time, the management team may convince themselves that every problem can be addressed, or that each problem has a low probability of occurrence. However, the collected problems may still be too much to manage. By its very nature, risk is statistical. It is possible to examine the collection of risks and make some projections about the project's likely success or failure. The result can even suggest that certain projects should be cancelled very early. Such projects can be rescoped and rebudgeted in a way that improves the focus and likelihood of success.

Figure 1: *Standish Group Project Results*

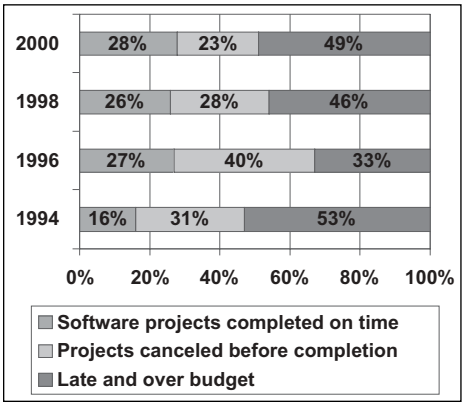

Risk Management Process

The Project Management Body of Knowledge (PMBOK) [2] includes a chapter on risk management. It describes the process steps as follows:

1. Risk Planning.
2. Risk Identification.
3. Qualitative Risk Analysis.
4. Quantitative Risk Analysis.
5. Risk Response Planning.
6. Risk Monitoring and Control.

"A standard definition of risk is an uncertain event that would cause an uncertain impact on project schedule, cost, or quality."

The metric proposed in this article fits the Qualitative Risk Analysis stage so it can be used as early as possible throughout the project duration. Rough estimates are available at this step, and are sufficient for an assessment of the overall project risk. However, the rough estimates will not suffice for risk items requiring real risk-response strategies such as mitigation and avoidance plans where more detailed work is needed.

In this metric, the distinction between steps three and four of the process model is important. The metric supports the viewpoint of senior management who wants to determine which of several projects has significant uncertainty. The project itself must deal with specific risks and quantitative analysis. As such, a risk manager on a large project will not find this metric as useful. He or she must have much more specific information.

History and Metric Definition

Risk is both old and new. The written history of risk begins in 1491 with the "Pacioli Puzzle," which arises from gambling when the game is stopped before completion [3]. The problem was solved by Pascal and Fermat in 1654 and so began the use of risk in forecasting. Today, risk is the core concept in insurance and has become a major focus in project management.

A standard definition of risk is an uncertain event that would cause an uncertain impact on project schedule, cost, or quality. Both the event and the impact have the element of uncertainty. The definition from probability theory is a bit more restrictive but it provides us with the metric:

R = P x V

The metric value of risk (R) is the product of the probability (P) of the event with the most likely value (V) of the outcome. If the risks are independent, we can add these estimates together for a combined estimate. So overall project risk is the sum of the separate risks.

Total Risk = Sum of all (P x V)

The Total Risk value and trends of Total Risk provide a picture of the project, making it easy for people to see some good and bad project patterns without delving into the statistical theory. The assumption about independence is necessary for the theory. However, in practice, risk management experts are aware that risks are not always independent. The metric is based on the theory derived from gambling where the assumption holds true.

Getting the Probability

There have been a few sociological studies showing the range of errors people demonstrate in estimating risk. Choosing an appropriate range helps when no historical data is available. Table 1 and its heuristics have been useful in avoiding the problems of underestimating and overestimating risk. Remember, most project managers see only three to five projects in their

career at any one company so they work from a very restricted sample. They need heuristics for estimates.

Five levels of probability seem to work well. Colleagues have not had a problem assigning an event to one of the recommended levels, so the suggested ranges provide good separation.

Label	Description	Value
Very Low	In your career, you have never seen this happen, but it could.	5% Range 1-9%
Low	It has happened on occasion.	25% Range 10-29%
Moderate	Sometimes it happens and sometimes not.	50% Range 30-69%
High	Most of the time this event will occur.	75% Range 70-89%
Very High	It has happened on every project, and you think it always will, but there is a chance of escape.	95% Range 90-99%

Table 1: *Risk Event Probability Estimates*

Analyzing the Impact

The impact of a project risk-event needs to be similarly divided into a few classifications and assigned a numeric value to manage risk. Making the numbers match conceptually when one risk affects schedule, another cost, and another quality or scope can be a bit of a stretch so a method is required to normalize the numbers.

A quick simplifying assumption works for the qualitative analysis stage: assign a single impact type. Choose from one of the following three image types: schedule, cost, or customers (sales). It is true that a risk event may affect more than one of these, however, coming up with a value for all the possible effects is challenging and probably not a worthwhile exercise until quantitative analysis. Narrowing the discussion of a risk event to a single type impact also focuses attention on the most useful response plans. This approach helps to avoid the problem of overthinking the impact of a risk. Here is an example of the kind of thinking to avoid at this early stage.

> Some employees are due for sabbatical leaves of two months. One may take that sabbatical during the project. You propose that turnover is a risk for the team. If this risk event occurs, it may cost some additional schedule time and additional resource cost to hire and train staff. If you lose schedule time, you may also lose some sales. What is the appropriate impact for this event – schedule, cost, or sales?

Experienced risk managers will understand that additional impacts will have to be considered when developing the risk-response plans.

Normalizing Risk Impact

The next challenge is normalizing the various impacts to arrive at a single numeric value for schedule, cost, and sales. Capers Jones reported that in 1996 "the typical project is 100 percent over budget when it is cancelled" [4]. This suggests that a useful normalizing factor is to set maximum risk impact at project cancellation. That impact value should be cost or schedule overrun of 100 percent, or when there is no customer or no potential first-year sale.

Of course, no project will be allowed to overrun to such an extent without senior management intervention, but that is precisely the point. Senior management should intervene when the uncertainty suggests the project is in trouble. Since the metric is applied at qualitative analysis, there is time to recover.

A Second Aside

> Why would we develop a product without customers? No one plans a project for a non-existent market, but the market can disappear or be misjudged. It happens all the time. Some well-known examples are the Newton tablet computer, the Iridium satellite telephone, and New Coke. Everyone also has an example of the *pet project* that was developed but was never used. Many organizations are surprised to learn that it is possible to cancel a project when sales or number of customers are factored into the risk management effort.

Using the possibility of cancellation as the highest risk impact, assign a value of five to cancellation. Five levels of risk should be enough. Creating the other levels again requires a bit of psychology. The PMBOK states an order of magnitude estimate is plus or minus (±) 35 percent of the base estimate. Using a range of 1 ± 0.35 is a range of 0.65 to 1.35. The ratio of these two numbers is 1.35/0.65 = 2.08, approximately a factor of two.

Thus to have a range that clearly separates the estimates, we must use a larger value. Using ±50 percent yields a ratio of 1.5/0.5, which equals a factor of three.

Experience suggests the psychology works, and people are comfortable with the results. Therefore, assign five to cancellation and divide the cancellation level by three successively to arrive at the other values. The following example points the way.

Consider a project that is scheduled for 18 months with a projected cost of $30 million and projected first-year sales of $27 million. This would be a project of about 100 people with about $5 million in external expenses. A risk event with an impact level of five would cause the following:

- Overrun by 18 months.
- Overspend by $30 million.
- Achieve no first-year revenue.

A risk event with an impact level of four (divide by three) would cause the following:

- Overrun by six months.
- Overspend by $10 million.
- Lose $9 million in sales (achieves $18 million).

A risk event with an impact level of three would cause the following:

- Overrun by two months.
- Overspend by $3.3 million.
- Lose $3 million in sales.

A risk event with an impact level of two would cause the following:

- Overrun by three weeks.
- Overspend by $1.1 million.
- Lose $1 million in sales (one customer).

A risk event with an impact level of one would cause the following:

- Overrun by one week.
- Overspend by $300,000.
- Lose $300,000 in sales (customer delays six months).

A useful interpretation is to say that the project manager can manage one or two risk events of impact level one within the project contingency and without unusual reporting. It would be necessary to generate a special report for any occurrence of impact level two. Any risk event at impact levels three or higher will require senior management's involvement to determine the response.

There is one more step in calculating the final impact. The numbers one through five were calculated by successive division by three. The final value has to put that back into a geometric scale.

Impact = 3^(level-1)

Impact Level	Impact Value
5	81
4	27
3	9
2	3
1	1

Table 2: *Impact Value Adjustment*

So a risk event with an impact level of five has an impact value of 3x3x3x3=81, as shown in Table 2.

The factor *three* is not arbitrary but is derived from the observation that order-of-magnitude estimates use a factor of two for the error range.

There is a temptation to turn the numbers back into dollars. This is a lot of work as revenue dollars are not the same as cost dollars or schedule days. The extra work makes sense for the top risks but not in general. Using the impact number instead of a dollar value also normalizes the risk metric across projects.

Risk Calculation

The final risk calculation follows the original equation:

Risk = Probability x Impact Value

The highest risk is 95% x 81 = 76.95
The lowest is 5% x 1 = 0.05

Normal usage is the sum of the highest 20 project risks. It seems that 20 risks at a time is a sufficient number to track for all but some mega-projects (over three years and more than 500 people). Barry Boehm, TRW professor of software engineering at the University of Southern California and author of the COCOMO estimating model, has suggested that projects manage the top 10 risks. There are two reasons this metric recommends watching the top 20. The first is that 20 risks x 5% = 100%. That is the recommended cancellation level so it makes for a convenient metric. The second reason is to make certain the project team investigates more than the first 10 risks to be certain that it manages the top 10.

Project Risk Score =
Sum (highest 20 project risks)

Implications

The Project Risk Score should be charted so senior management can see scale and trends. Since there is a threshold (threatened cancellation) implicit in a risk with impact level five, that threshold should also appear on the chart. An impact level equal to five translates into an impact value of 81. Figure 2 is a sample chart from an actual project.

There are many implications in the chart and its use. The threshold is a powerful concept. Senior management will focus a lot of attention on a project that is above the threshold. The fact is, projects with risk higher than the threshold simply will be late, over cost, or fail to meet project quality goals. Some projects have risk levels that are astronomically high. It is theoretically possible to see a value of 1,539 with 20 risks that are very likely to occur and have an impact rating of five. Of course, such a project should be cancelled and restarted. I have actually seen only one project risk value over 400. That project had to make major changes to deliver even a subset of the desired functionality. If the threshold concept had been introduced at the start of that project, it would never have gotten into so many problems.

A somewhat opposite situation also can occur when a project shows particularly low risk. The project manager or senior management may have a sense that a project is at significant risk, but the metric does not show it. Use that low number as a signal that a risk collection effort is needed. The project manager must gather a wider audience and run a facilitated session to identify those other risks. Make sure to include stakeholders from other locations and groups outside the development team. Develop the organization taxonomy for risks like the one in the "Continuous Risk Management Guidebook" [5] from the Software Engineering Institute to make the data collection more complete and rigorous.

A normal response when the project risk is high is to manage that risk down. This can happen several ways:

- The time for the risk event may pass without incurring the problem.
- The team may adopt an avoidance plan so that the event cannot occur.
- The team may adopt a mitigation plan to reduce the impact.
- The team may transfer the risk to another organization.
- The event may occur and the project eats the contingency.

The last four responses cause the project to incur a specific cost that should appear in the project planning and reporting. Each of the responses requires the project manager to make some update to the risk database.

Finally, product managers (not project managers) should be hesitant to select a project of very low risk. If the risk is so low, why not address a more aggressive product plan? Risk avoidance is not generally a winning strategy in the marketplace. The point is to manage risk to appropriate levels for the organization, product, and project. Risk management is a systemic study and not a technological one.

Figure 2: *Sample Project Risk Score*

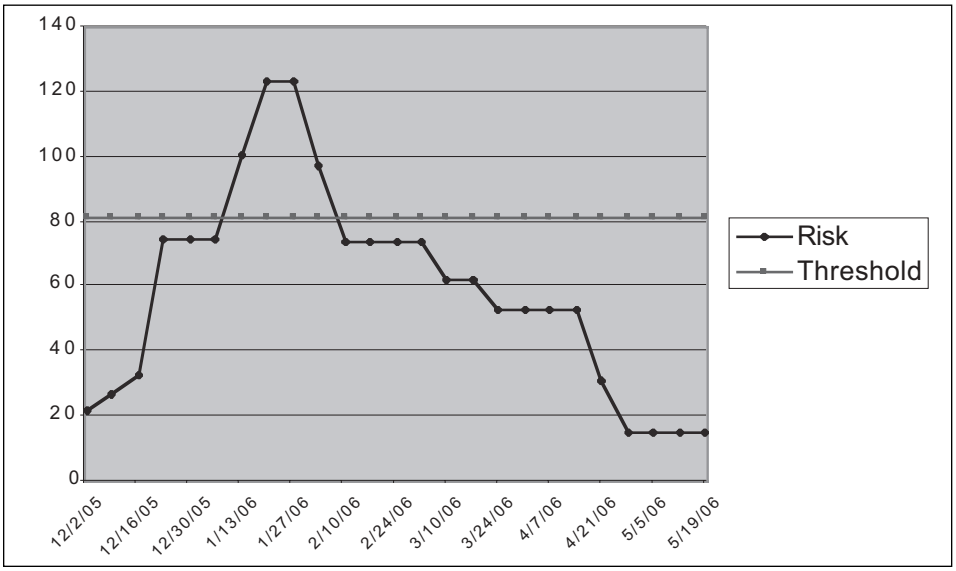

Implementation

There are several challenges in adopting the project risk metric. The following is a list of the top challenges:

- **A database for collecting and managing risks.** There are a number of products that will do the job. Implementing one will require the addition of project and sub-project identification and organizational process support. This work cannot be institutionalized without an automated system.
- **A process model.** The basic framework is available in the PMBOK, the Continuous Risk Management Guidebook, or the Institute of Electrical and Electronic Engineers standard for risk management [6]. The process model has to be extended to cover a risk taxonomy that is appropriate to

the organization.

- **Automated reporting.** Chances of success are better if the project risk chart is automated and is required as a part of the regular project management review. The risk metric should be checked at least monthly.
- **Training.** Training is a big effort. Training project managers to do risk management takes days, not hours. Writing good risk scenarios requires at least eight hours of training and much practice. Learning the organization taxonomy of risk takes time. Evaluating impacts probably takes three hours of training. Directors and senior managers also need at least three hours of training. Do not attempt to implement a project risk metric without decent training on risk management.

Summary

Many seasoned project managers say that advanced project management is mostly risk management. This metric makes that statement visible and concrete to a much larger audience. It provides fast visibility and has a high emotional impact on managers.

The project risk metric, however, has been tested in only one location and on only a dozen projects. The simplifying assumptions made in order to develop and use the metric make it suspect for use by risk practitioners who must perform detailed quantitative analyses and develop risk mitigation and avoidance plans.

It does provide a comparison between projects that is useful to senior management. If senior management is presented with one risk at a time, they are likely to develop a confidence that they can deal with each risk as it comes. Dealing with each risk separately and successfully may convince them that the project cannot really be in trouble. Management may then come to believe that the project team is whining about problems instead of dealing with problems, and real risks may not be addressed in a timely fashion. Presenting senior management with a picture of the total project risk will encourage them to take appropriate systemic actions when these are necessary. Product managers on projects with high risk will need additional justification and resources to add scope. The development team may have an easier time getting training or adding consultants when needed.

The key is presenting senior management with better visibility into the project so that project change-management becomes faster and easier, and finally, so that product delivery becomes predictable.◆

References

1. The Standish Group International, Inc. CHAOS Chronicles Ver. III. West Yarmouth, MA: The Standish Group, 2003 <www.standishgroup.com/chaos/toc.php>.
2. Project Management Institute. A Guide to the Project Management Body of Knowledge (PMBOK Guide). Newton Square, PA: Project Management Institute, 1996 <www.pmibookstore.org/productdetail.asp?productid= 4106>.
3. Bernstein, Peter L. Against the Gods: The Remarkable Story of Risk. Hoboken, NJ: John Wiley and Sons, 31 Aug. 1998 <www.wiley.com/Wiley CDA/WileyTitle/productCd-0471295639.html>.
4. Jones, Capers. Patterns of Software Failure and Success. Boston, MA: International Thompson Computer Press, 1996.
5. Durofee, Audrey J., et al. Continuous Risk Management Guidebook. Pittsburgh, PA: Software Engineering Institute, 1996 <www.sei.cmu.edu/publications/books/other-books/crmguidebk.html>.
6. Institute of Electrical and Electronic Engineers. "Software Engineering: Software Life-Cycle Processes, Risk Management." Proposed Standard. New York: IEEE, 2004 <http://standards.ieee.org/announcements/pr_p1540. html>.

About the Author

Robert W. Ferguson is a member of the Technical Staff at the Software Engineering Institute. He has more than 30 years of software development and management experience in several industries. Ferguson is a member of the Computer Society of the Institute of Electrical and Electronic Engineers and the Project Management Institute. He has been active in the software process improvement community for several years and is past chairman of the Chicago Software Process Improvement Network.

Software Engineering Institute
Carnegie Mellon University
4500 Fifth AVE
Pittsburgh, PA 15213
Phone: (412) 268-9750