

Managing Consulting Project Risk

PRESTON G. SMITH CMC, GUY M. MERRITT

Have you ever experienced significant surprises toward the end of a consulting engagement, when you had little time to recover and few options for responding? Think about these surprises for a moment. Did you or someone in the client organization have a premonition that such a “surprise” might occur, because it had occurred before with your other clients or in their experience with other consultants or projects?

Project risk management is a process of identifying potential surprises up front and managing them throughout the project to diminish their likelihood or impact. You will not be able to avoid every potential pitfall that might threaten your project, and some are truly unknowable in advance. But with explicit attention to project risk, you can improve your odds considerably.

The key to managing risk is being proactive about it. As just suggested, many risks are half-known in advance, but often they are dismissed in the rush to get on with the “real” work of the project. Or perhaps there is some denial involved. Risk management is simply a matter of acting explicitly in advance to prevent a risk or diminish the risk’s consequences. Being proactive is easier for some individuals and organizations than for others. The opposite of being proactive is the “fire fighting” behavior that many managers enjoy and some organizations reward tacitly. If your client organization embraces fire fighting—dealing with problems only when they be-

come crises—proactive risk management will be more challenging to apply.

As business becomes more competitive, the pressure tends to drive us toward ignoring *possible* problems and thus toward more reactive behavior. However, many organizations today are recognizing the high cost of dealing with project problems that could have been anticipated, so project risk management has become a popular management topic.

Understanding Risk

If you ask ten people to define *risk*, you will get ten quite different answers. With such divergence of opinion, you may have difficulty achieving the consensus needed to be proactive about managing risk. Moreover, most of the material on risk management, such as what you may have learned in an MBA program or what you will find if you search the Internet, is not helpful for managing project risk. Financial risk (insurance, investment portfolios), for example, provides few tools for dealing with a risk proactively. Medical risk is another large area, but it is similarly limited in managing projects.

To be manageable, a risk requires three characteristics:

■ *Uncertainty*. A risk is a potential event, and you do not know in advance if it will actually occur. Therefore, the probability of occurrence

■
Identify consulting project risks and manage them, lest they wreak havoc down the road.

of a risk is always less than 100%. If an event is certain, then we instead call it an *issue* and take action on it differently than for a risk.

■ *Loss.* A risk always has the potential for causing a loss, which could be measured in financial terms, time, corporate image, follow-on consulting engagements, or other terms. *Potential* is a key word here, because, if the risk does not occur, there is no loss. In managing project risks, we ignore the possibility that things could turn out even better than expected—not to be pessimistic, but because our objective is to minimize the loss.

■ *A time component.* For a risk to be manageable, it must have a limited time frame. For a consulting project, many risks will end with the termination of the engagement. But some risks will end sooner, and some—such as nondisclosure agreements—may live for a period beyond the engagement. The time component could be expressed as a condition that determines when the risk ends, rather than directly in terms of time. For example, if you are conducting training as part of an engagement, a risk might be that not enough people sign up for the training. This risk terminates when you reach your minimum class size.

A Model of Risk


A risk can be dissected into components that guide us in managing it, such as those shown in the risk model in Figure 1. The starting point of this model is the *risk event*—a crisp statement of the risk that concerns you. For example, if your concern is not obtaining sufficient participants for a class you are planning, the risk event might be stated as “Insufficient attendees to hold the class.”

The next component is the *impact* of the risk—the loss that could result if the risk event occurred. In the training example, the impact might be the project delay due to rescheduling and promoting the class, or depending on your frame of interest, it could be a personal loss of so many billable days.

To the right of this in Figure 1 is the *total loss*—a number that represents what would be lost if the risk event and its impact occurred. We prefer to express total loss in ei-

Figure 1 A MODEL OF RISK

This model shows components of the risk essential to managing it.



Source: Adapted from Fastrak Training Inc. training material. Used with permission. © 1996.

ther time or money, such as the number of days lost to rescheduling the class (time) or the personal cost in lost billings (money).

At the bottom of the figure are two critical ingredients: *risk event drivers*—facts in the project environment that cause you to believe that the risk event will occur, and *impact drivers*—facts in the project environment that cause you to believe that the impact will occur.

Risk event drivers for the training example could be (1) it will be summer vacation time, and (2) people dislike the older training facility being used. If we presume the rescheduling rather than the lost billings impact, corresponding impact drivers could be (1) five days will be needed to reschedule the facility and reissue the announcement, and (2) participants must be given three weeks to sign up.

This may seem like overkill, adding unneeded complexity to the situation. However, the drivers perform two functions essential to managing a risk effectively. One is to quantify the risk in order to evaluate its importance and compare it against other risks. You can see already from the impact drivers that your project will slip a month if you must reschedule the class. Such facts associated with the drivers move risk management from the realm of guessing to one of analysis. Even more important, the facts associated with the drivers are crucial to for-

mulating effective mitigation plans. For example, when you identify vacation time as a risk event driver, you can look for ways to resolve the risk.

Going back to Figure 1, the two remaining boxes at the top are *probability of risk event* and *probability of impact*. Multiplied together and then multiplied by the total loss, these two quantities yield an *expected loss*. Expected loss is an overall measure of the risk—the average loss you can expect from it. For example, if you conducted many of these classes with this client over a period of time, sometimes you would get sufficient participants (no loss), and sometimes you would have to repeat the offering (total loss). The expected loss tells you how much delay you can expect on average.

This model, with a risk event and an impact in tandem, is our preferred model. You can also use a simpler one, with the risk event and impact combined, or you can add even more stages to the model. As far as we know, none of the existing project risk management literature uses any kind of a risk model explicitly. Many use such models tacitly, often implying a simple, single-stage model. We believe that explicitly using a model greatly aids in thinking about a risk and mitigating it effectively. Otherwise, conflicting opinions within a group tend to block the consensus needed to mitigate project risks proactively. Retaining two stages in the model becomes important when we turn to planning responses, because risk event drivers lead to a different kind of response plans than impact drivers do.


The Risk Management Process

Figure 2 illustrates a five-step project risk management process. In contrast with risk models, which appear to be unique to our approach, most of the literature on project risk management suggests a process similar to the one illustrated here.

Such a process will have to be scaled to the size of the project at hand. For short engagements, one or two people might complete the whole process in less than an hour; but for large projects involving many consultants, it might take several people a week.

Figure 2 FIVE BASIC STEPS OF PROJECT RISK MANAGEMENT

The first four steps are usually done once, but the last one is ongoing.



Risk Identification

This is essentially a brainstorming process to uncover any risks that could potentially afflict your project. You will not pursue many of these very far, but the few important ones may require considerable attention, even to the point of convincing you to terminate or radically revise the project.

Brainstorm using any technique you prefer, but always provide a clear initial problem statement. Beyond this, do not judge contributions now. Encourage piggybacking to take one idea in another direction, and push for off-the-wall ideas that pertain to the problem statement. Go for quantity now, as this breeds quality in the end. Brainstorm with a group of up to a dozen for large projects, or as few as one or two individuals for small ones. Try to get at least one person from the client organization involved, as their viewpoint, experience, and objectives will

PRESTON G. SMITH CMC

consults to manufacturers on product development techniques, especially for rapid development. He is coauthor of *Developing Products in Half the Time*.

(preston@NewProductDynamics.com, www.NewProductDynamics.com)

GUY M. MERRITT is a staff program manager working for a major communications equipment vendor. He has specialized in developing and deploying risk management processes on various research and development projects.

(gmerritt60544@yahoo.com)

be quite different from yours. In general, you want to bring a breadth of experience to bear on uncovering potential risks.

Depending on the project documentation available and the nature of the project, a variety of thought starters can prompt risk discovery. In general, risk in a project lies in the interface areas, so focus on the interfaces between the consultant and the client, between departments of the client organization, between phases or tasks of a client process, or between geographic areas. The project schedule—especially if it clearly shows dependencies between tasks—can aid in pinpointing risky areas. Or use a process map—either of your consulting process or of a client process that you are investigating—that shows interfaces between organizations.

Another way to uncover risks is to first paint a picture of project success. Then ask what could get in the way of achieving this picture. If you conduct similar engagements repeatedly, a longer-term approach is to build a prompt list of problems you have experienced on past projects, and then use it to suggest potential problems on the current project. Lastly, ask the client directly what could go wrong. When we conduct client interviews, one question we often ask client managers is, “What concerns do you have about the success of this project?”

As you identify each risk, word its risk event and its resulting impact crisply. If you are working with a group of people, you can put these on sticky notes that you post on the schedule, process map, or whatever chart you are using to prompt your risks, one risk per note.

At this point, you have finished the most interesting, obvious part of risk management. Unfortunately, the process often ends here, which is worse than not starting it in the first place. Not only do you receive no benefit from preventing risks if you fail to pursue them from here, but also you set yourself up for frustration later when some of the risks you have identified start happening. Conclusion: Do not even identify your risks unless you intend to do something about them.

Risk Analysis

Look back at Figure 1. For each risk, you now have its risk event and its impact. The objective of risk analysis is to fill in the rest of

the boxes on the chart so that you can estimate the overall magnitude of the risk, its expected loss.

Typically, you start by determining the risk event drivers, and then move on to the impact drivers. If there are no facts leading you to believe that the risk event could occur—or that its impact could occur if the risk event occurs—then you do not have a risk. Thus, many risks will drop from your list at this point. Clearly, you will have to decide how strict you will be about facts. If a certain problem occurred on eight out of ten of your last projects, for example, you might consider it a fact.

We cannot overemphasize the importance of listing your drivers. Everything from here on depends on these. Not only will you use them to estimate how serious a risk is, but they will become the foundation for mitigating it—if the drivers cause you to decide that it is serious enough to mitigate.

Next, determine the total loss that you or the client could suffer if the risk event and its impact occurred. As was the case for risk in the training example, there may be several possible impacts. Pick the most serious among them, and express total loss in monetary terms, time, or some other metric. In the next step, risk prioritization, you will be comparing risks based on their corresponding loss, so it will be convenient to have them all based on the same quantity.

Last, determine the two probabilities. These should stem from your drivers. On the surface, the most appealing way to describe probabilities is to just use qualitative terms—such as low, medium, and high. This may be adequate for simple projects, provided that you establish the meaning of these terms (anchor them). However, using such so-called ordinal scales is fraught with theoretical problems; for instance, you cannot perform arithmetic with such values. Consequently, we use numbers for probabilities, and to eliminate endless bickering about whether a probability is 43% or 47%, we allow only certain values, such as 10%, 30%, 50%, 70%, and 90%.

Explicitly using a risk model greatly aids in thinking about a risk and mitigating it effectively.

A data management tool such as a spreadsheet is useful for tracking these data. Use one row for each risk, and start with columns for each box in the risk model. Add an initial column that identifies each risk uniquely, such as R1, R2, and so on.

Risk Prioritization

As mentioned earlier, calculate the expected loss of each risk by multiplying the two probabilities together and then multiplying that amount by the total loss. If you have been successful at quantifying all risks on the same basis, initial prioritization is then simply a matter of using the spreadsheet software to sort your risks by their expected loss. If you used a qualitative scale or have expected losses expressed in varying units, such as time *and* money, there are alternatives. For instance, you can sort the risks manually by placing a risk that ranks high on both total loss and probabilities at the top of your priority list.

You can also use a risk map, which does the same thing graphically while avoiding arithmetic by separately displaying the total loss of a risk on the horizontal axis and its combined probability on the vertical axis. Any risk in the upper-right corner of such a map warrants attention, and one in the lower-left can probably be ignored for now.

You may decide to treat some risks in a special way. For example, catastrophic risks have very high consequences (total loss) but may be so unlikely to occur that their expected loss is not near the top of your list. However, their consequences would be so horrendous that you may decide to manage them anyway.

Prioritizing your risks is an important step, because from here on you will be investing resources in the risks you choose to manage actively—resources that must be diverted from doing the “real” work of the project, that is, actually completing the project deliverables. You will not have enough resources to manage all the identified risks, so choose to manage only those that pose the greatest danger to your project.

Risk Planning

From prioritization, you now have a short list

of the risks you wish to manage actively. With a few exceptions, each risk will receive one or more action plans. You have many choices of action plans:

- *Acceptance*. Simply decide that you will accept the risk, that is, do nothing about it. This can occur when its consequences (total loss) are small or if you discover that action plans would be more costly than if the risk event occurred.
- *Avoidance*. Strange as it may seem, once you understand a risk, there are sometimes ways of using another route to avoid it.
- *Transference*. Transfer the risk to another party, such as a contractor or supplier.
- *Redundancy*. Simultaneously pursue a parallel route, deciding at some future date which of the two routes to take. This is usually an expensive option.

This leaves us with the most common types of action plans:

- *Prevention plans*. These seek to change the risk event drivers in a favorable way. Prevention plans are truly proactive, since they attempt to keep the risk event from occurring.
- *Contingency plans*. These seek to change the impact drivers in a favorable way. Contingency plans are more reactive, because they diminish the consequences only if the risk event actually occurs.

For action plans to work, they must be taken seriously. This means they become another task in the project and receive a budget, schedule, and labor resources—no differently than any other project task. If your action plans receive second-class treatment, they are unlikely to resolve your risks.

Risk Monitoring

In contrast with the previous steps, you repeat this one regularly throughout the project. Monitor the following items:

- Progress on action plans for actively managed risks.
- Removal from the actively managed list of those risks that have been mitigated effectively.

■ Addition to the actively managed list of any risks being monitored that have become more serious and now merit active management.

Also conduct a mini-version of the risk identification session to scan for any new risks that may appear as circumstances change. If you discover such a risk, pass it through the other steps, as necessary.

How often do you conduct this ongoing monitoring? As often as you normally consider the project budget or schedule. The objective of managing project risks is usually to preclude schedule and budget problems. If you do not give attention to risk management as frequently and seriously as you monitor the schedule and budget, which are more reactive measures, then you simply are not acting proactively.

An Example

To illustrate the process and how the model is used, we provide an example from a consulting project.

Suppose that your consulting engagement is to assess how well the client organization is employing a certain software package that, among other things, is intended to tie together their global operations uniformly. You will assess the organization's capability from structured interviews of users and from associated data collection and analysis. From this assessment, you may recommend training, software upgrades, improved help desk capability, better linkages to other corporate systems, and similar improvements. Refer to Figures 1 and 2 as we proceed through this example.

Among other project risks, the risk identification step reveals a concern that several interviewees in the Stuttgart office will not be available when you have scheduled the European interviews, necessitating a return to that office to complete the interviews. In this example, we focus on this one risk, stated as: *Insufficient interviewees available in Stuttgart on the days that I schedule interviews there.* The resulting impact statement is: *I must schedule a second trip to Stuttgart.*

The associated total loss is \$10,000 (a

couple of billable days plus travel expenses for another trip to Germany). Observe that there are alternatives to this impact. For example, you may decide, if the risk event occurs, to forego the Stuttgart interviews and suffer a loss of assessment quality instead. Likewise, you may decide to express the total loss as three weeks of project time lost rather than \$10,000, if time is more precious than money for this engagement. The point is that the risk model helps you to think through this situation carefully and proactively.

Next, you list your risk event drivers (refer to Figure 1):

- RE1. Stuttgart is a sales office.
- RE2. Salespeople travel extensively.
- RE3. Stuttgart interviews are scheduled for December (holiday period).
- RE4. Stuttgart personnel have little interest in this assessment.

Then list your impact drivers:

- I1. I will be visiting London and Milan offices just before and after Stuttgart, respectively.
- I2. Corporate global sales meeting is in San Diego in January.
- I3. Hamburg plant has a human resources department with some interviewing capability.

These two lists of drivers will help you to set the respective probabilities at the top of Figure 1, and they can also help you establish the magnitude of the total loss. More important, they will suggest prevention plans to reduce the likelihood that you will fall short with your Stuttgart interviews (your risk event). For example, you might be able to shift the Stuttgart interviews to November or January (RE3), and you could publicize the importance of your project in Stuttgart (RE4). If these prevention plans fail and your Stuttgart interviews fall short, your impact

Do not even identify your risks unless you intend to do something about them.

drivers will suggest contingency plans to avert another trip to Stuttgart:

- Rearrange the London and Milan visits, or insert some free days in this week to provide flexibility (I1).
- Catch some of the missed Stuttgart interviewees when they are in San Diego (I2).
- While in Europe, train someone in the Hamburg HR office to conduct the missing Stuttgart interviews (I3).

All of these possibilities would probably occur to you when the failure in the Stuttgart interviews was upon you. But by proactively identifying and dealing with this risk early in the project, you have many more options for resolving it. For instance, when you share with your client that they may have to pay \$10,000 for a second trip to Stuttgart or suffer the consequences of this hole in your assessment, they can become your partner in mitigating the risk. A proactive approach on this risk will also diminish the ancillary risk that the \$10,000 will come out of your own pocket.

Implementation Guidance

These risk management techniques offer significant benefits—mainly in making projects more predictable—but they can have significant costs, too. Consequently, we recommend modifying them to obtain the greatest benefit for the least cost. The best way of doing this is to start with a somewhat more formal and complete process than you anticipate you will need. Then conduct a project retrospective to see what is working and what is not, and adjust the process to provide greatest advantage.

The terminology of project risk management is a great source of confusion. This is why we have used terms carefully and consistently. As you spread the techniques with clients, be prepared to train them in this method. Otherwise, you may spend much time arguing over risks and drivers or educating client personnel one by one.

Initially, we mentioned that the earmark of good project risk management is its proactiveness, and we observed that this opposes the reactive, fire-fighting style of many managers. To the extent that you or your clients gravitate toward fire fighting, risk management could pay great dividends, but it will also require a substantial change in organizational behavior.

Project risk management is not likely to last long if it is merely an appendage to the project. Therefore, to be successful, integrate it into the project seamlessly; for example, as a normal part of schedule, budget, and project meetings.

Incorporating risk management into your consulting projects will likely require an investment of time and effort, even if you scale it down as just suggested. We have found, however, that what it really does is shift the way you spend your time on a project. Rather than working after-the-fact to overcome problems, you spend your time up front in keeping many of them from happening. The net effect is that the outcome of your project is more predictable. ■

Note

1. This process and the underlying risk model are described in detail in *Proactive Risk Management: Controlling Uncertainty in Product Development*, by Preston G. Smith and Guy M. Merritt (Productivity Press, 2002).