

What Are the Chances? Overcoming Barriers in Assessing Risk Probabilities

Tom Kendrick, Program Director, UC Berkeley Extension, Berkeley, CA

Abstract

Project risk assessment relies on estimates of both risk impact (consequences) and risk probability (likelihood). Estimating probabilities tends to be the more difficult of these two, for a number of reasons. Because probability estimates are often so underestimated, many significant risks may be overlooked.

Estimating risk probabilities more realistically results in improved risk management, lower overall stress, and more successful projects (to say nothing of earlier intervention to adjust or abandon project concepts that are prone to failure). This paper explores barriers to probability estimation and describes a process for determining and refining estimates. It also discusses how to use better probability estimates to assess overall project risk, and how to use ideas such as ROI simulation and “Value at Risk” (VaR) to validate project financial assumptions. The paper builds on concepts outlined in Kendrick’s *Identifying and Managing Project Risk*, Second Edition (AMACOM 2009, recipient of the 2010 PMI Cleland Literature Award).

The Problem

Estimating is hard

All estimating is difficult. Most project managers will readily admit to a lack of expertise, even when estimating concrete quantities such as duration, effort, or money. Probability is abstract, and even more problematic. Because people lack much intuition about probabilities, we do not understand them well. For most of human history, we did not even bother to study them. From the time of the Greek philosophers well into the Renaissance, people assumed that the odds in gambling and other situations involving chance were controlled “by the gods” and therefore beyond our comprehension. (This is well covered in Peter Bernstein’s excellent review of the history of risk, *Against the Gods*.) Today, despite the fact our theoretical understanding of statistics and probability has evolved significantly, most people are still bad at assessing probabilities, as evidenced by the many lost “bar bets” and questionable decisions we make. And while this can cause problems in many areas of life, it can be especially dangerous for projects.

“Unlikely” risks tend to be overlooked

Because the probabilities estimated for many project risks are inaccurately low, many potentially significant risks may be excluded from a risk register because they appear too unlikely to worry about. Even those that are listed may be disregarded or marginalized during assessment because of their assumed minuscule likelihood. Ignoring risks does not remove them from a project; it converts them from potentially controllable project aspects into potentially rude, sometimes disastrous, surprises. Dealing better with risk begins with understanding why probability estimates are so often wrong.

Key Challenges to Risk Probability Estimating

Listing significant risks

Developing a good sense of project risk, including realistic probability estimates, starts with a robust list of project “uncertainties that matter.” Making risks and potential problems visible is most effective when the risk identification process is integrated into the processes of project definition and planning. As project information accumulates, good project leaders explore assumptions and planned work for what is missing, incomplete, or otherwise not adequately understood. They ask about worst cases and what might go wrong or interfere with the deliverables, schedules,

budgets, and other components of their project documentation, listing all exposures and uncertainties encountered during project analysis and planning, without regard to how likely or unlikely they may appear at the time. This results in a much more robust and complete list of project risks than a perfunctory brainstorming session conducted with team members in the final stages of planning, when everyone is anxious to stop planning and start working. Starting with a realistic list can make a huge difference in effective risk management.

Too little relevant information

The probability for any specific project risk will always be somewhere between zero (no chance of occurrence) and one (inevitable occurrence). Risk probabilities must all fall within this range, but picking a value between zero and one for a given risk poses difficulties. There are only three ways to estimate a probability. For some situations, such as flipping coins and throwing dice, you can construct a **mathematical model** and calculate an expected probability. In other situations, a simple model may not exist, but there may be many historical events that are very similar. In these cases **statistical analysis of empirical data** may be used to estimate probabilities. Such actuarial analysis as this is the foundation of the insurance industry. In all other cases, probability estimates are based on **guessing**.

Projects are complicated, so developing a computational model is generally either impossible or requires more effort than it appears to be worth. Data relevant to events that may seldom if ever occur is also sparse for most projects, so statistical analysis can rarely be used for estimating risk probabilities. As a result, project probability estimates usually rely on guessing, based on analogous situations, scenario analysis, “gut feel,” and wishful thinking. For too many project risks, a lack of information results in imprecise probability estimates.

Plausibility versus probability

When things are uncertain, people devise scenarios based on what they consider to be plausible options and alternatives. It’s human nature to focus on what we would prefer to happen, so the scenarios we create tend to emphasize the most desirable outcomes. We think in stories, and when starting from a vacuum we invent one—often the story that we most prefer. The plausibility bar for such creative scenario building can be quite low. When considering likelihood, if there are more visible options that avoid unpleasantness, our estimates will tend to favor them and underweight the more adverse alternatives. People consistently tend to confuse plausibility with probability (or, worse, implausibility with impossibility).

Optimism and biases

Optimism is essential to successful project management. We must believe that the projects we work on are possible and that we will prevail. Excessive optimism, however, conceals what we need to know and undermines our chances of success. Because we “hope for the best,” desirable outcomes tend to seem more likely than is realistic. (“The Cubs are sure to win the World Series this year—they’re due!”) Conversely, people estimate undesirable outcomes as less likely. (“No way it’s going to rain on our outdoor wedding.”) This optimism bias is even more pronounced when dealing with small probabilities. Most people perceive little difference between one chance in a hundred compared with a chance in a million. Both seem “impossible.”

Errors in estimating bias are compounded by heuristics—“shortcuts” we all use for analysis and decision making. The effect of **anchoring** can be substantial in any estimating process. Giving excessive weight to unimportant (or even random) data can interfere with any process generating numerical forecasts. Hearing or thinking about a number will bias the result toward that number. A single overly optimistic person (or project leader) can anchor the whole team and cause chronic probability underestimation. Because of the **availability** heuristic, we tend to overemphasize easily remembered or recent data, ignoring older information and unpleasant experiences that we are trying to forget. Another potential source of bias is the **representativeness** heuristic. When we lack directly relevant information, we plug in something that we do have or can easily get. Sometimes the information is sufficiently related and can help answer otherwise difficult questions, but we do this even when the data is completely irrelevant to the case at hand. Instead of trying to estimate the probability, for example, we might be tempted to poll the project team to see how many people think a risk might happen. The percentage of contributors who find a given risk plausible can be interesting, but it is not a good way to ascertain probability. There are many other potential pitfalls related to heuristics; these are just a few of the most significant. A good in-depth exploration of this can be found in *Thinking, Fast and Slow* by Daniel Kahneman.

Bias does not always contribute to inappropriately optimistic probability estimating. Anchoring can also potentially skew probability estimates to the high side. If there have been recent significant problems, the availability heuristic

may result in overestimates. Although such cases may help to counter excessive optimism, they too are sources of estimating inaccuracy that you should monitor and manage.

Lack of statistical expertise


Finally, even people who have studied statistics (including people who teach it) find explaining probabilities and understanding them daunting. For example, descriptions of medical studies are significantly easier for people to understand when they use numerical case counts (concrete data) rather than percentage-based presentation (abstract data) of exactly the same information. This becomes even more confusing when dealing with “percentages of percentages.” Few project staff members have very much education in or experience dealing with statistical concepts.

A Process for More Accurate Probability Estimating


Probability 101

Much of statistics is complicated, using lots of difficult mathematics. The parts that relate to project analysis, though, need not be terribly convoluted. The primary concept is the mean, the center of a distribution. In most cases the simple arithmetic mean of a data population is sufficient for project discussions. In discussing risk, the standard deviation, a measure of the expected (or measured) variability about the mean, is also useful. When considering a collection of risks, incorporating correlation information into the overall project risk analysis provides additional insight. While it is possible to introduce a wide spectrum of other statistical ephemera into project analysis, however, most important questions can be addressed with a solid understanding of means and standard deviations.


That said, there is a lot of variation to be seen when considering different types of data populations. Much of the information surrounding projects is made up of concrete, measurable quantities (duration, effort, money, etc.). These quantities tend to lend themselves to modeling using a “typical,” distribution having a single peak in the middle and diminishing tails extending to both higher and lower values. The Gaussian (normal) “bell-shaped” distribution is of this type and so are other commonly used models such as the triangular distribution. (See the figures below.)



There are many, many other statistical distribution types (Beta is a particular favorite), but those usually applied have a central peak tend to result in comparable analysis. When dealing with after-the-fact empirical data sampled from data populations, there are many examples that tend to roughly resemble the peak-with-tails shape of the theoretical models when graphed as a histogram, as in the case below:



Any project leader can look at a histogram such as this, understand the calculated mean and standard deviation, and easily interpret what it conveys. However, probability data (as with most data related to binary—rather than continuous—outcomes) does not look like this. In retrospect, all probabilities are always either zero or one (or one hundred, as a percentage). When plotted, such data looks like this:



The mean of 10% does make some sense, indicating that historically there is about one chance in ten that the type of risk measured here will occur. Also, the standard deviation is quite large because all the data is crowded at the extremes of the possible range. What is apparent is that the structure of the underlying sample data varies from situation to situation. This affects what the statistics mean and has implications for how best to use them.

In any event, an understanding of at least some basic statistics and the underlying “shape” of the data is essential to useful probability estimates.

Team inputs

If more than a few project contributors are in a position to help with probability estimates, get them involved. Make sure that estimates are individually created by each knowledgeable team member working solo to avoid “groupthink,” anchoring, and clustering. When soliciting risk probability estimates, encourage people to use “what if?” and “worst case” thinking to balance any optimism bias. Gather the estimates provided and use Delphi analysis to segment them into thirds: highest, lowest, and intermediate. Explore the reasoning behind the highest estimates. Probe for specifics from those who provide probability estimates that appear to be too low. Work to counter “wishful thinking.”

History

Even though relevant archived data may be sparse, investigate what you do have. Probe for both data and anecdotes. Review project reports, and retrospective “lessons learned” analyses from earlier projects. Discuss potential risks with peers, stakeholders, and team members. When speaking with other people, do it “off the record”; what actually happened in some situations may differ from what was written down.

Nonlinear assessment scale

Initial, **qualitative** probability assessment depends on ranges. Setting ranges for which people have little intuitive feel can result in meaningless or misleading data. When setting ranges for probability estimates, linear scales such as 20% quintiles are very difficult to work with. Selecting a specific risk probability “bucket” from among the middle three quintiles inevitably relies mostly on guessing, and gaining meaningful consensus can involve a good deal of unproductive debate.

Ranges for which people have some personal connection work better, and such scales tend to be nonlinear. When “highly likely” corresponds to 50% or higher, it can be aligned with personal experiences flipping coins (or any other common binary process with equal odds). “Moderately likely” can be about 20%, and compared with rolling a specific number using a fair six-sided die. For low likelihood events, with a probability in the low single digits, a comparison can be made with rolling a double six with a pair of fair dice, drawing a card needed for a winning poker hand, or with some other well-understood unlikely event. (The more unlikely the comparison event becomes, however, the less accurate our intuition becomes—hence the willingness novice poker players have to draw to an inside straight.)

Nonlinear relative scales with more ranges can be devised. You could construct a five-level scale (with breakpoints such as 0%, 4%, 10%, 20%, 40%, and 100%) for assessing relative risk probability. While this will probably be easier for people to use, unless you can align the breakpoints with experiences that people can relate to it may not actually improve the accuracy of your analysis. Qualitative probability assessment is never easy, but using scales that people understand and not including too many ranges will make it a lot less confusing and therefore more realistic for prioritizing key risks.

Similarly, **quantitative** probability assessment must be “relatable.” The best quantitative probability estimates align with earlier qualitative range estimates, and generally fall near the high end of the qualitative range.

Skepticism

Being too credulous gets project leaders in a lot of trouble. Successful risk management depends on **asking a lot of questions**, especially concerning aspects of the project that appear dubious. For probability estimates that seem too low, ask questions to explore what the assessment is based on.

One effective way to reframe the estimating process and detect issues is to introduce a monetary element. Ask people if they would wager a hundred dollars, using their probability estimate as odds on the bet. If the probability estimate is 50%, you would stand to lose \$100 if the risk happens and win \$100 if it does not. If the estimate were 10%, the odds would be 9:1; you would lose \$900 if the risk occurs and win \$100 if not. Loss aversion can provide a substantial counterbalance to excessive optimism—even when the losses are hypothetical.

Refining the Estimates

Outliers

Be very skeptical of “outlier” risks that have probability estimates that are significantly lower than others or than would seem to be reasonable. For these risks, discuss the specific consequences of occurrence, including measurable criteria (money, schedule slippage, overtime, etc.) and less tangible factors such as damaged reputations, future business prospects, reduced trust, stress, and interpersonal conflict. Ask the owners of the work related to the risk to reevaluate the probability estimates after documenting recovery efforts, costs, and other consequences.

Bayesian techniques

Probability estimates should not exist in a vacuum. Better estimates result from conscious use of historical data and trends. Bayesian analysis provides concrete guidance for doing this, based on work that originated with Thomas Bayes in the 18th century. But the concept of using prior experience to explicitly adjust future forecasts remains alien to many people. Examples of this include believing that losing sports teams that are “due,” and treating common problems as though they are rare exceptions. Bayes’s theorem provides a way to incorporate actual historical data into probability estimates for upcoming events.

Whether or not you wish to formally apply Bayesian techniques, it’s useful to spend some time considering whether your probability estimates are consistent with trends associated with comparable situations. If your probability estimates seem too low (not justified by improving trends in your archived data), revise them upward.

Evaluation of optimism and bias

Optimistic bias is inherent in many probability estimates. Examine estimates, especially those that appear “off,” for evidence of common sources of bias, such as anchoring, availability, and inappropriate representativeness.

Discussing “positive risk” in projects has become popular in some circles in the past few years, even though examples of significant *uncertain* opportunities are somewhat rare. One good use of this idea, however, is in reframing “threat risks.” Instead of looking at the probability of a threat in a given baseline plan, we can consider the opposite situation: we assume the risk will occur in the project and treat the possibility of avoiding it as an opportunity. As an example, assume that a specific project risk—such as getting a key component late—has an expected impact of ten workdays of schedule slippage and an associated probability estimate of 15%. To test whether this estimate makes sense, consider a plan that incorporates the risk impact as a given, by increasing the relevant duration estimate(s) by a total of 10 workdays. In this project, there would be a “positive risk” of reducing the schedule by 10 workdays (by receiving the part earlier), having an associated probability estimate that is the complement of the original percentage, or 85% (100% - 15%). If assuming an 85% probability that you can shorten the project by 10 workdays because the component will arrive “early” seems unreasonable, adjust the probability to be more realistic, and reflect this adjustment in the estimate for the original risk.

Uncertainty about uncertainty

Rarely are probability estimates definitive numbers for which you can have high confidence, for the many reasons already discussed. Impact assessment estimates are similarly imprecise for many risks, which compounds the issue. When you cannot determine the consequences of a risk with much accuracy, assess the potential range of possible outcomes. Discuss the underlying causes of both the mildest and most severe impact that could occur, and reexamine the probability estimates for each extreme. Update your probability estimate for the risk either using “error bars” based on the range information you develop. Alternatively, you can adopt a percentage estimate toward the upper end of the range.

Integrating the Overall Risk Analysis

Reserves

One of the most common uses of risk probability and impact estimates for the overall project is to justify and establish management reserves. Using realistic probability estimates and worst-case impact data, review the adequacy of schedule reserves and budget contingencies (if any). If inadequate (or no) reserves are set, use your analysis to build a case for realistic reserves based on risk analysis.

Expected impact

Even when the probability of a risk is realistically estimated to be very low, it’s never zero. Determine whether the worst-case consequences for a given risk can be tolerated for your project (or your organization). If the potential

results of a risk, regardless of how unlikely it may seem, exceed what you can afford to bear, plan to respond to it. In extreme cases you might even reconsider the project as a whole, and either modify or abort it.

When the best probability estimates you have are uncertain, base your risk prioritization and response strategies primarily on expected impact. When you are in a situation involving dire consequences, it's not unlike the scene in the film *Dirty Harry*, where Clint Eastwood's character says, holding his (possibly empty) gun to the head of a miscreant, "Well, to tell you the truth, in all this excitement I kind of lost track myself. But being this is a .44 Magnum, the most powerful handgun in the world, and would blow your head clean off, you've got to ask yourself one question: Do I feel lucky? Well, do ya, punk?"

Facing project-threatening consequences, it's unwise to feel lucky.

Simulation

Many projects, particularly large, complex projects with significant risks, use schedule simulations to explore the impact of timing risks on the deadline and other key project dates. This is useful, but for risky projects with substantial budgets, financial return on investment (ROI) simulations can be even more illuminating.

Much of the project risk data having financial impact involves increased costs associated with project activities and deliverables. Additional uncertainties arise from the estimates of expected benefits and any uncertain opportunities inherent in the project. Basic business cases are generally based on single-point estimates of expected cost (which tend to be underestimated) and single-point estimates of anticipated revenues and/or benefits (which are usually wildly overestimated). The simplest ROI assessment is the difference between these two numbers—using neither any risk information nor adjustment to compensate for the time value of money—and provides for most projects a breathtakingly rosy picture with which to justify the project. Incorporating risk into the picture provides a more realistic view.

As an example, consider a project with expected financial benefits \$1,000,000, and a cost of \$750,000. The ROI based on this is \$250,000, or roughly 33%. But this is probably not the whole story. There are doubtless both risks and some uncertainty around the benefits, so a more realistic view might look like the data in this table:


	Assumed Values	Assumed Probabilities	Expected Values
Project Cost:	\$750,000	100%	\$750,000
Project Benefits:	\$1,000,000	100%	\$1,000,000
Risk 1:	-\$500,000	10%	-\$50,000
Risk 2:	-\$200,000	15%	-\$30,000
Risk 3:	-\$250,000	10%	-\$25,000
Opportunity 1:	\$50,000	5%	\$2,500
Opportunity 2:	\$125,000	25%	\$31,250
Total Expected Value			\$178,750

Assuming the probabilities are realistic, the expected return falls a bit, but it remains at about 24%—not bad. Neither risks nor opportunities ever “partially” occur, however. As discussed earlier, they either happen or they don't. Given this view, the possibilities for this project are summarized in the following:

Nominal Value (Certainties Only):	\$250,000
Expected Value:	\$178,750
Maximum Value:	\$425,000
Minimum Value:	-\$700,000

The maximum looks like very good news; if all the good stuff happens and none of the bad, the return is excellent. If the reverse happens, though—we see all of the risks with none of the upside—the result is disastrous. The expenses more than wipe out the return and the project shows a significant loss.

Neither of these extremes is very likely, though. For a more nuanced view of this, computer simulation can help. Using the data in the table and doing 1,000 trials, the average return is about \$170,400, which is only slightly less than the “expected” return calculated using the figures for risks and opportunities weighted using their estimated probabilities. The values can be plotted in a cumulative graph, to look like:



This looks a lot more interesting. Almost 20% of the time, this project will either return nothing or lose money. About 30% of the time, its return will not exceed the expected value in the table above. However, nearly half the time (49.3%), the project will return \$250,000, and there is about one chance in six it will do even better than that. Given this, especially considering the potential for a loss, does this still look like a good project?

Of course, even this view, eye-opening as it is, is not entirely complete. The possibility that the project will fail or be cancelled is not included (there ought to be a probability estimate associated with the benefit number along with all the other uncertainties). In this case, there will be significant out-of-pocket expenses up to the point when the project ends, and no, or at least very little, benefit delivered.

Value at Risk

Value at Risk, or VaR analysis has become very popular in recent years. VaR is a technique for estimating the maximum amount of loss that can be expected from a financial investment. It came into widespread use in the 1990s and built on the work in portfolio theory done by Harry Markowitz. It incorporates risk analysis into ROI calculations, and attempts to show the level of risk being assumed so it can be better managed. VaR is based on some reasonable, but not completely bullet-proof, assumptions. When financial markets are well behaved, VaR analysis allows financial firms to eke out a slightly larger return while remaining at the level of risk they believe

appropriate. When the markets are volatile, though, as they were in the global 2007–2008 financial meltdown, VaR fails. VaR in those conditions misrepresented the risk being assumed, and it contributed substantially to the debacle.

The fundamental idea of VaR is to determine probability distributions for the investments and time frames under consideration. Using the distributions and computer analysis, potential returns or losses may be calculated. VaR is stated in terms such as “\$100 million for one week at a 95% confidence interval” This means that there is no more than a 5% chance expected of losing in excess of \$100 million dollars in the next week. An excellent description of the history and mechanics of VaR can be found in the book *Strategic Risk Taking* by Aswan Damodarian.

Getting VaR to work requires selecting and applying an appropriate probability distribution. Determining the distribution is done using some combination of the same three techniques used to estimate probabilities: Mathematical modeling, Empirical analysis of historical data, or guessing. As in most complex situations, historical analysis and guessing tend to dominate, but even plugging in a simple model such as the Gaussian distribution may provide useful insight when time frames are very short and the rising and falling variations are small and for the most part in balance.

VaR can fail, though, for a number of reasons:

- A selected probability distribution is a forecast, not a guarantee. Actual results may (and often will) vary.
- Even if the probability distribution has generally correct parameters, the “shape” may be inappropriate.
- Data used to define the distribution may be incomplete or otherwise inappropriate.
- Conditions assumed to be stable may prove to be drifting or otherwise more volatile than expected.
- Other assumptions may be unwarranted due to defective analysis or dishonesty.

Applying VaR to a project requires examining a longer time horizon than for many other investments. The time scale extends out to the breakeven point for the project, which can be months or even years, so “project VaR” involves investments that move in slow motion compared with VaR’s usual application. The VaR objective, however, is the same—assessing how much money is at risk with our project investments. The analysis for a given project begins with an ROI-type analysis similar to that in the preceding example. Because the time horizon for a project is long, using a typical assumed distribution and a high confidence interval, VaR would be quite large compared to the project budget (unless risks are few and there is close to zero probability of cancellation). For the example project, assuming that the information is realistic and substantially complete, there appears a 95% chance of losing no more than about \$250,000.

All of the potential problems cited above apply to project VaR. For one thing, there is a general lack of relevant data and bias in estimates for probabilities and for probability distribution parameters, as explored earlier in this paper. This, if coupled with a less-than-completely-honest and thorough analysis, will undermine creation of useful VaR assessments.

Nonetheless, at the organizational level, VaR for a portfolio of projects, especially projects having adequate risk assessment and probability estimates can provide a useful way to assess and manage financial exposure from project investments. One challenge with this, as with all VaR portfolio analysis, is to properly model correlations among projects. Most projects in organizations tend to be highly correlated, which results in increased overall risk and higher VaR.

Summary

Better probability estimates for project risks result in greater understanding, improved plans, and more successful projects. To achieve them:

- Be thorough in identifying project risks. Use history (data and anecdotes), comprehensive analysis, and brainstorming, and integrate risk identification throughout all project planning.
- Resist confusing plausibility with probability.
- Understand sources of bias, and work to minimize them.
- Err on the high side for probability estimates.

- Use “probability/impact” risk data to evaluate and validate project business cases and expected returns.
- Be skeptical. Be **very skeptical**.

When the thought “What are the chances . . .” occurs to you, the chances are likely higher than you think.

References

Bernstein, Peter (1998). *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, Hoboken, NJ.

Damodarian, Aswan (2008). *Strategic Risk Taking*. Wharton School Publishing, Upper Saddle River, NJ.

Kahneman, Daniel (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, NY.

Kendrick, Tom (2009). *Identifying and Managing Project Risk*, Second Edition. AMACOM, New York, NY.

Kendrick, Tom (2012). *Results Without Authority*, Second Edition. AMACOM, New York, NY.

Project Management Institute (2008). *A Guide to the Project Management Body of Knowledge (PMBOK® Guide)* (4th ed.). Newtown Square, PA: Project Management Institute.