

Lean Risk Management

Rick Bollinger, MSE, PMP
Senior Consultant, Natural Systems Process Improvement, Inc.

Abstract

Lean risk management is a combination of tight estimating, optimal buffers and squeezing more of what can be ascertained from project current and historical information. It includes close monitoring of risks, quick triggering of contingencies, and integration of risk management with project planning and regular project activities.

A comprehensive system of risk management is described that avoids a statistical problem and calculates risk buffers that protect project objectives with estimated degrees of confidence. Common textbook methods for estimating buffers are based on shaky assumptions, and these are examined and improved upon.

Risk probabilities and impacts can vary over time. Methods are shown to integrate risk management into project schedules so that estimated costs and dates accurately reflect changes during the entire lifetime of the project. The role of risk management in maintaining a portfolio of projects is described.

A strong case for disciplined risk management is made and suggestions and techniques presented to make the process more lean. Lean in the senses of efficiency, eliminating waste, improving processes and creating value for customers and the organization.

The Top 10 Reasons We Don't Do Risk Management

With respect to David Letterman, let me lay out ten reasons we do not perform risk management.

10. We are all OPTIMISTS at heart.

Although experience teaches us differently, we ignore these lessons and our optimism takes over. Plus, injecting pessimism into planning always inflates the numbers. This is not welcomed by upper management or customers.

9. Hope is ALWAYS a strategy.

The times that hope has proved a fruitless strategy never seem to matter. Every time seems unique. Every failure has a new reason. If we do not recognize the pattern, we are only left with hope. It is the default strategy.

8. We all remember some lucky project.

Reasons for having HOPE may be many. Failures are so common and varied that they are not memorable. What is really memorable is some project from our past where everything worked. We take credit for its success. Unfortunately, luck probably had a larger role. However, that is forgotten in the halo of good feelings that remain. We are waiting for it to happen again.

7. Don't lessons learned fix everything?

It is unlikely that proper accounting of things done right or wrong is done in the aftermath of panic and blame (or euphoria and back-slapping) following project terminations. This is made more difficult when the repository of accumulated wisdom is on a network drive that is not backed up, or is as poorly organized as a junkyard.

6. Are you feeling lucky today? (Punk!)

Some people feel that they must prove their worth by taking chances. But, not everybody is as well prepared as Dirty Harry.

5. Who wants to display uncertainty/ignorance?

Quite often risk management is not attempted out of fear. It may be new territory. Even with training, it can be daunting to start in a new environment. A project manager may feel insecure trying to develop this new skill.

4. Spouting whales get the spear!

More apt than saying the squeaky wheel gets the grease. Speaking badly about the chances of a project may get a person uninvited to the party. Who wants to work next to a downer?

3. Risk management is really DEPRESSING.

Contemplating how a project might fail could maybe jinx it? Certainly it does not buoy the spirits. This is finger pointing in the future tense. People might need therapy afterwards.

2. We have no DATA.

Starting risk management is difficult. Estimates of risk probabilities and impacts are just guesses without hard data. Unless projects record their successes and failures in risk management, there will be no data for future projects to use. Somebody has to jump in and swim first, no matter how cold the water is.

1. This is PLANNING and we need to get to WORK!

Upper management can be very intolerant of risk management. Plans wait while a black cloud passes over the project. The new estimates of cost and duration cause bad feelings all around. The hopes and dreams and careers of stakeholders are dashed on the rocks of uncertainty that are revealed. Meanwhile, the unfinished schedule slips.

Top Reasons for Doing Risk Management

Identify potential problems

It is healthy for a project to think about the ways it can get into trouble. Getting blindsided by something that could easily have been predicted is embarrassing. Often prevention is simple or cheap. This contributes **value** to the project. Do your homework. It might impress somebody.

Handle them proactively

Risks and problems are like fires. When they are small, they are easier to handle. When they are big they are dangerous. Early intervention is the most **efficient** strategy. Position fire extinguishers where they can be easily reached. Mitigate future impacts by building preventive mechanisms or constructing contingency plans and facilities. Put all these actions in the plans with their effort and costs. It is better than insurance. Or, finger-crossing.

Be their masters, not their slaves

Gain mastery over risks by having plans to mitigate them directly and build resources to act if they occur. Make it so they cannot surprise you. So they cannot kill you. So they cannot control you. So you control them.

Keep them in mind

The price of freedom is eternal vigilance. Or, something like that. Don't let a risk sneak up on you and **waste** resources. Don't let the problem fester unseen. Pick your head up and look around every now and then. What is new? What has changed? What actions need to be triggered?

Learn from them

Risks and problems are great teachers. What does not kill you can make you stronger. How well were mitigation and contingency actions? Are there opportunities for **improvement**? Save the experience in the organization's memory.

Protect project objectives

Properly formed risk mitigation and contingency planning can be targeted towards specific project objectives. Build buffers and walls. Have multiple paths to the goals. Preserve **value** for stakeholders. Plan victory like a military campaign.

Cover our butts!

Imagine Dilbert with his manager. Dilbert states an 'estimate'. But, to his manager it is received as a 'promise'. This is true for many of us. Giving out an early estimate that does not reflect the true consequences of risks and uncertainties can be dangerous to your health. They will make a project more expensive and take longer. So, they should be included. Also, never bake an opportunity into an estimate. Always hedge your bets. A low rough estimate will stick in the mind of an upper manager or stakeholder. You will be punished when you try to take it back.

OK, let's switch gears. Let's talk about risk management in the context of some root concepts of lean: **efficiency**, eliminating **waste**, process **improvement**, and creation of **value**. We can begin with fixing a current problem, and then move to ways to get better project performance (and profits!) from risk management.

Using Classic Buffer Calculations

There is a problem with the typical method of calculating buffers in all the popular risk management books. Let me explain. It starts OK with a concept called 'exposure'. This is the product (e) of the probability (p) and the impact (r) of each risk (i).

$$e_i = r_i * p_i$$

Many books then say to add up the exposures and get a total exposure for the project (E) or phase, et cetera. Then, set your buffer (B) to this number. Supposedly, this amount of buffer is all the protection a project needs from its risks.

$$B = E = \sum_n e_i = \sum_n r_i * p_i$$

A book titled *Waltzing with Bears* even says that by "...setting budget and schedule reserve equal to budget and schedule exposure, you are allocating a reserve that is sufficient, *on average*, to contain your risks." (DeMarco, 2003, p 69, emphasis added.) I will demonstrate below that this is false.

What everybody also seems to have missed is that this calculation is the same as the one for the average or mean R (R_{bar}). And, we know that an average is in the middle. Thus, roughly half of the values will be above the mean, and half will be below. So, using this method, **HALF OF YOUR PROJECTS WILL BE LATE AND/OR OVER BUDGET** if we are just considering risks. (See Exhibit 1, below.)

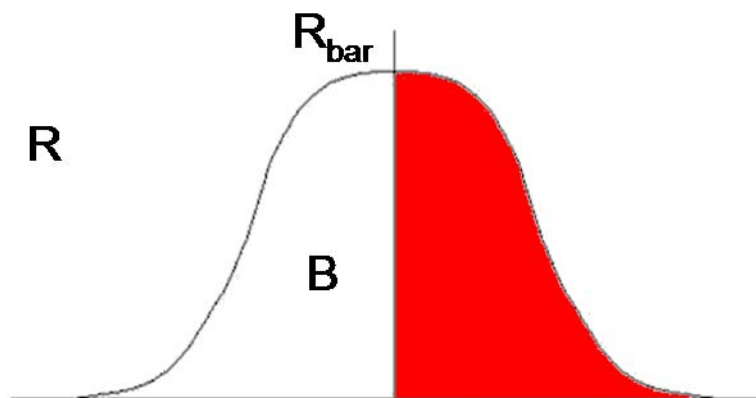


Exhibit 1 – Example distribution of outcomes about a mean value

This calculation for B accounts for the MEAN value, but NOT for the standard deviation. This is throwing information away. This is WASTE. We can do better.

Risks can be described by random variables. Their outcomes will have variation about the mean value. Classic buffer calculations do not give any indication about how much to the left or right of the mean to expect the outcomes. As project manager I don't care very much about the performance of reserves "on average". I care about the performance of reserves on the current project.

A BIG Assumption That Almost Never Holds

Waltzing with Bears is wrong because it is based on a BIG assumption that almost never holds. The assumption is that the value (cost) of resources gained or lost before the deadline are the same. This situation is diagrammed in Exhibit 2, below. IN REALITY THIS IS HARDLY EVER TRUE.

In the fortunate situation when the project finishes early, (before the vertical line for the buffer value labeled 'B'), then the resources can be moved to other projects. People can be reassigned; unspent dollars and hours can be reallocated on a one-for-one basis without penalty. This is a good situation, but it cannot be made to happen every time.




Exhibit 2 - Equivalent resource values before and after deadline

In the much more likely situation, we know that if you go over budgeted costs or schedule, then there are typically penalties to be paid, by circumstance or according to contract. If the deadline is missed, then the hours worked must be paid at overtime rates, if that is the approach taken, or maybe worked using hours or resources stolen from other projects. This can put those other projects at higher risk as well.

The real situation looks more like Exhibit 3, below. Beyond the deadline, resources become much more valuable and cost much more to purchase. In addition, the credibility of the personnel and company are burdened with the failure. (The little upturn at the left end reflects the possibility a project may be audited if it comes in much earlier or less expensive than expected.)




Exhibit 3 - Realistic resource values before and after deadline

Thus, if contingency buffers and risk reserves are calculated as many textbooks recommend, you might not be protecting your projects as well as you think. You are betting with a confidence level of 50%. It is no wonder that experienced project managers fudge. Fudging way more than you need is WASTE.

Expressed this way, I imagine your customers would not sponsor a project that had a 50-50 chance of failing. Certainly, most project managers would not accept such a project. *Waltzing with Bears* (DeMarco, 2003, p 69) also mentions that a "...more defensive strategy would be to allocate something more than aggregate exposure, while a less defensive strategy would be to allocate less." I am in favor of fudging, but on a rational basis.

How This is Based on Rational Choices

If you are not convinced yet how this BIG assumption is false, let me demonstrate with an example from almost everyone's experience. Have you ever been driving in a construction zone and been exasperated that there was no



Exhibit 4 - A Rational Choice?

activity? I mean, you are sitting in your car in the sun, with orange cones and barricades all around you. All traffic is squeezed into one lane. You are getting more late for work or a meeting. And, for no reason you can figure there are no workers anywhere to be seen. The equipment, if any, is sitting still. (Like in Exhibit 4, to the left.)

I know the reason. It is because of risk management. The contracting company has taken the resources off this project and put them towards other projects. A typical situation that calls for this is that one or more of the projects are in trouble, or trending bad, and at risk of missing deadlines.

At play are big bonuses or big penalties. It has been decided that the project in which you sit is to be sacrificed. The penalties that would be paid by the other projects are too high. The resources that are freed

up can be best allocated to make sure that other projects succeed. The bonuses or profits from those contracts will make up for the losses or penalties on this project. They are gaming the system, and you are the victim.

They are only making rational choices. And, so should you.

How Big Should Buffers Be?

Obviously, risk buffers need to be larger, but just how much larger is a good question. They need to be at least as much as the exposure. We don't want to bet AGAINST the project. And, we would want a statistical basis for the new buffers. We would want the new buffers to increase the confidence levels we would have on our estimates. And, we would want, for business reasons, for the buffers to be EFFICIENT: no bigger than are needed. The better an organization is able to estimate, the closer it can bid to the amount that wins the bid and still makes money.

Central Limit Theorem to the Rescue

I never thought I would ever use the Central Limit Theorem in statistics. But, it has great value here. It says the sums of random variables tend to become approximately normal, i.e. they follow the Gaussian Curve. And, this applies even if the underlying random variables are not Gaussian. (See Exhibit 5, to the right.)

This is fortunate because we are exactly interested in the sums of risks. And, statistically, we know a lot about the normal distribution. In particular, we can estimate its variance and standard deviation. So we can get some idea of the spread (or dispersion) of risk outcomes.




Exhibit 5 - Gaussian Curve

Maybe we can figure out a useful fudge factor that could really protect projects and their objectives.

Estimating the standard deviation of the curve allows us to make reasoned decisions about how much buffer we need to give us confidence that we will get the protection we need. If we can get some idea of the underlying risk distributions, then, we can estimate the standard deviation of the resulting sum.

Let's get started by assuming that our risks are independent of each other. This is usually true. In case some risks are not, then they have the same probability of happening. Just add their impacts together and treat them as one. A group of independent risks can be represented by a collection of random variables with the Bernoulli (n=1) distribution. They either happen or they don't. We can estimate the probability and impact of each risk, and then derive an estimated standard deviation of the outcomes.

The variance for each risk is then $r^2 p(1-p)$ where r is the impact of the risk and p is the probability. So, we can derive an estimate of the variance for the sum total. Since variance is a second degree operator and we have no constant multiplier:

$$VAR(\sum X_i) = \sum VAR(X_i) = \sum (r_i^2 p_i(1-p_i))$$

So the standard deviation to expect of the outcomes, added all together, where r_i is the impact of risk(i), and p_i is the probability of risk(i), is the square root of this:

$$S. D. = \sqrt{\sum (r_i^2 p_i(1-p_i))}$$

Using this we can arrive at a statistically relevant buffer. In Exhibit 6, below, you can see a representation of a project. It shows a starting point and a duration (D) that is estimated without considering risks. The sum of risks has an additional contribution with a resulting distribution to the right of D . An estimate can be chosen based on the degree of protection desired by selecting a multiplier, z . Multiplying the standard deviation by zero results in the standard textbook buffer. As you can see, doing that there is plenty of red that is not covered. As higher values of z are applied, more protection is provided. Also, more cost and/or time is implied in the estimates.




Exhibit 6 - Estimating Total Duration with Risk Buffer

Choosing Your Buffer

As you can see in Exhibit 7, below, $z = 0.0$ gives you the same result as the classical buffer calculation. $z = 1.0$ gives you about an 84% confidence in the estimate. $z = 2.0$ will give you 98% confidence. (Only 2.28% would be expected to exceed their buffer.) One cannot arbitrarily pick a higher z . Estimates have consequences. Note that the additional buffer is not necessarily a large portion of the total. R can be much less than D .

z value	0.0	0.5	1.0	1.5	2.0	2.5	3.0	3.5
% expected to be over	50%	30.85%	15.87%	6.68%	2.28%	0.62%	0.13%	0.02%
% expected to be under	50%	69.15%	84.13%	93.32%	97.72%	99.38%	99.87%	99.98%

Exhibit 7 - Choose your 'z'

Using simulation we can compare the protection of robust buffers to that of classic buffers. Exhibit 8, below, shows the result of one iteration of an Excel model constructed. It models 10 independent risks with impacts ranging from 100 to 1000. The actual value is chosen randomly in each iteration of the model. The probabilities are also random from 0 to 100%. Exactly 100 samples are run each time. The simulation rolls the dice, so to speak, for each sample and adds up the damages for each risk that is manifested.




Exhibit 8 - Buffer Comparison

The top bar shows the total of risk impacts in the results of one such iteration. The second bar shows a robust buffer calculation, and the bottom bar shows the classic buffer calculation. This is typical, but not always the case. The classic buffer is about half the total impacts and the robust calculation is between the classic calculation and the total. The larger the 'z' value used, the greater the portion the robust calculation will be of the total. The configuration of the bars above will not always hold. All the calculations use random values. But, the trend is clear and the way to bet is straightforward.

One must remember that the values above would be tacked onto the end of actuals and estimates of project cost and duration done without considering risks. How the buffers scale to these values depends entirely on the size of the risk impacts. Protection might be a very expensive component of project totals, or just a small part. What is important is the part that is played in the endgame. Adding appropriately to your confidence level will pay off handsomely as deadlines loom.

A caveat that must be observed is that we may find ourselves working with weak assumptions. Are the risks really independent? Are the number of risks sufficient for the sum to be modeled using the Gaussian distribution? These must be answered, not with mathematics, but with reality. Does the cost implied by a robust calculation seem outrageous? Is it too low? What CAN be said is that a robust calculation has as much basis as a classic calculation. It can do no worse, and in many cases it will do much better.

What is true is that we are now working with more information, not less. We have extracted more information out of the data we have. So, maybe, instead of worrying about the assumptions not under our control, we should try to improve the source data with which we are working: the estimates of risk probability and impact. We can derive

great **value** by keeping good records on the frequency of risks, their impacts, and the costs and successes of mitigation and contingency plans and activities. This we CAN control.

In fact, I would recommend that a risk management system be established. But, that is the subject of another presentation. In the interim, we can still operate in our improved situation. Let's take a look at what that would be like.

Structuring Projects for Risk

To reiterate, risk buffers are not typically a major portion of a project's total budget. The reason is that they are not based on the length of the project. They are based on the impact of the risks which are components of a project. However, there is reasoning that can be applied if the magnitudes of risks become very large.

Lots of BIG Risks

In this case the new buffer calculations are very important. Make sure all the non-independent risks have been combined. Pay close attention to the cost of each **z** that needs to be added. They will be significant. It will be important to find the point where risk and cost are balanced.

One BIG Risk

If there is one major risk of high impact, you might consider making the buffer 100% of the impact and plan for it to happen. This might not be acceptable to your customer, or it might be one that would invalidate the project. So, you might try making the whole deal provisional. That would require agreement to start the project knowing the risk and then abandon the project if the risk manifested. It is wise in this situation to establish sensitive triggers so that the plug can be pulled at the slightest hint of bad news. This will cut the losses of still operating before termination.

Using the New Buffers and Benefits

You might note that the **actual** durations and costs of projects will not increase using the new buffers. Those are a function of the actual **work** involved and the **problems** encountered (unless customers cancel).

We cannot affect the **randomness** of risks any more than before. They will happen or not happen as before, assuming the same mitigations. However, we can substantially improve the context in which they happen and the robustness of our preparation and responses to them.

Customers will not pay any more or wait any longer for project deliverables. But, expectations can be made more in line with the risks involved. Customers should actually find that you are **early** and **under-budget** according to the '**z**' factors that you choose. You will more seldom over-promise and less frequently under-deliver. You won't have to charge more. You will just have to apologize less.

Better risk management can contribute to more and better business. Accumulated knowledge of risk probabilities and impacts can be exploited to improve estimating. As confidence increases in the total estimates, customers will have increased confidence in your projects. In bidding situations, yours will be closer to the ideal amounts and with greater safety. Your competition will have to bid more than you and lose the contract, or less than you at greater hazard of being unprofitable. That creates **value** for your organization.

Implementing Lean Risk Management

Plan your strategy for risk management

A good way to start is to plan in the beginning how risk management will be done. Develop a starting set of risk sources and categories of impact. Start with some training on risk management. Start risk management for the first project and develop a template with which to manage the effort. Also create a template with which to record historical data.

Identify risks

Hold a meeting and brainstorm potential problems for the project. Assign risk owners to watch over them, develop mitigation and contingency plans, and represent them within the project.

Estimate risk impacts and likelihoods

Estimate probabilities and estimated impacts for each risk. Record these in a template. Spreadsheets work very well in this role.

Calculate risk reserves/buffers

Using the new equations, calculate buffers for the project and/or for specific project phases, if risks apply to them.

Show risks and plans in the schedule

Insert the buffers into the schedule and budget. They work in many tools like Microsoft Project®. Risk groups present much like tasks. They can have duration and cost. If a risk group applies to a specific phase, then create such a buffer task at the end of that phase and before the start of the following ones. Risks that apply to the whole project can be grouped and estimated for a Management Reserve at the end, or at the beginning. (See Exhibit 9, below.)

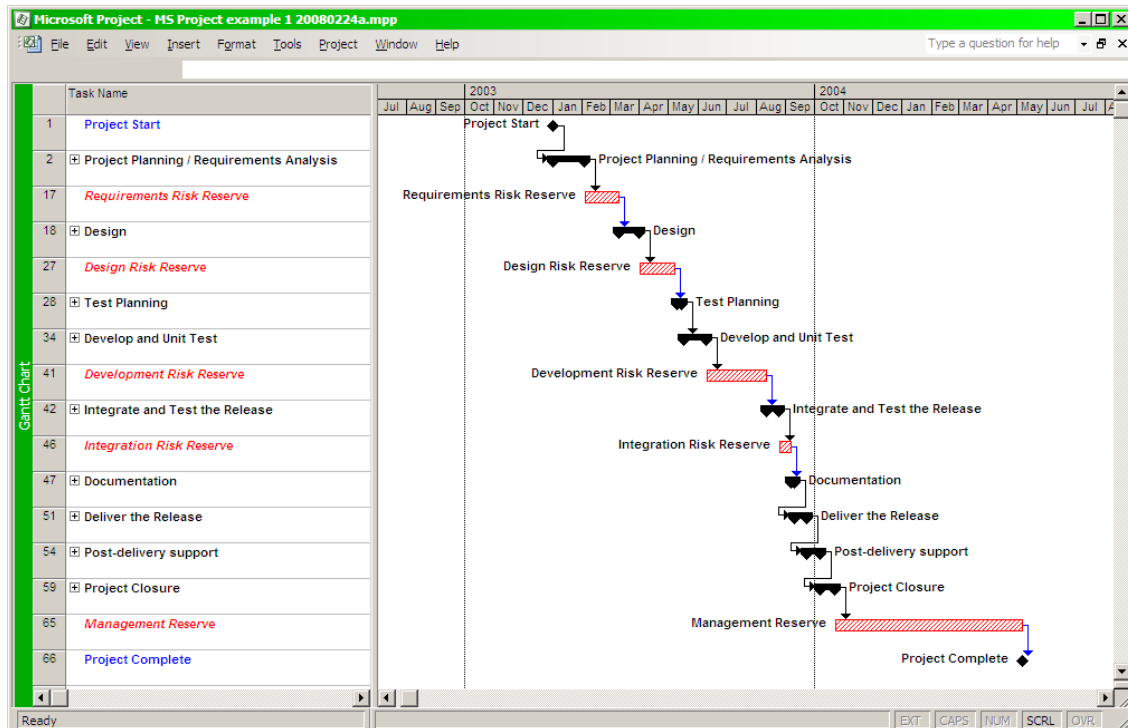


Exhibit 9 - Inserting risk buffers into a schedule

Not only buffers should be inserted into schedules. Mitigation plans that have tasks to execute should also be scheduled and worked on like other tasks. Their priority depends on how big and close their risks are. You can also build the tasks of your contingency plans into your schedule. Show them in parallel to the risks and isolate their costs. Scheduling tools can also be used to calculate the duration and costs involved. This is a handy way to estimate the impacts of a risk. And, when they manifest, now the project knows how to respond.

Monitor risks during execution

On a regular basis check that the probabilities and potential impacts of risks have not changed. Remember, as the project progresses, the opportunities for a risk may increase or diminish. Completion of a phase without

manifestation of one of its risks means it can be deleted. Near the end of a project, almost all of the risks should have been retired. And, if the opportunity for a risk has evaporated, then a proper accounting of the event would be a recalculation of the buffer to which it was allocated. This reduction would bring in a deadline or lower a budget. However, it is almost hardly ever in the nature of a project manager to do this.

Account for risks in schedule/budget

Good record keeping can be leveraged by future projects. So, it is a good idea to decide ahead of time the budget categories into which actual time and costs will be applied. The following guide (Exhibit 10, at the right) may get you started. For example, mitigation actuals could be applied to the mitigation tasks. This could be an accounting category for future projects. And, contingency actuals could be tracked against their risk reserve tasks.

Activity	Actuals Allocation
Mitigation planning	Planning
Contingency planning	Planning
Mitigation execution	Mitigation tasks planned
Contingency execution	Risk reserve 'tasks'
New risk response	Management Reserve

Exhibit 10 - Allocating risk activities

It is a good idea at the termination of a project to go back over risks and tally up their costs and successes. Also, confirm the organization's risk sources and categories. Do all the risks fit into the current model? Do changes need to be made?

Improving Practices

Finally, we improve risk management by learning more about the ones that afflict us or we are able to avoid. This enables us to estimate better and plan better. We also do better by increasing what we remember about problems and risks. Risks turn into problems, and past problems are future risks. We can apply lessons learned and use this information in future mitigation and contingency plans. All of which begs for disciplined risk management and a comprehensive system in which to catalog and calculate.

Another use of a risk management system is in support of managing a portfolio of projects. A risk management system can give us indications of the health of our projects early on. Like Earned Value Management, we can gauge the trajectory of a project and the certainty of making budgets and deadlines.

The new buffer calculation eliminates **waste** of information. The reasoned choice of 'z' values increases the **efficiency** of estimating and bidding. The implementation of better buffers for the project and phases creates protection for the **value** of project objectives. And, the accounting and recording of historical data forms the foundation for future **improvement** in risk management and estimating. These are aspects of making risk management lean.

Absent a systematic approach, we are operating ad hoc. Even still, it is possible to look at a portfolio of projects and risks to work out a strategic approach to their management. We might even find ourselves in the situation of terminating a project for its resources so as to optimize the portfolio to our advantage.



References

DeMarco, T., Lister, T. (2003) *Waltzing with Bears*. New York: Dorset House Publishing Co., Inc.