



RISK MANAGEMENT

Your Company Is Too Risk-Averse

by [Dan Lovallo](#) , [Tim Koller](#) , [Robert Uhlener](#) and [Daniel Kahneman](#)

From the March–April 2020 Issue

In theory, companies are supposed to create value for stakeholders by making risky investments. And as long as no single failure will sink the enterprise, those investments may be quite large. It won't matter if even a significant percentage of them fail so long as the success of other bets compensates, which usually happens. It's an approach to investment that's supported by economic theory going back to the 1950s work of Nobel laureate Harry Markowitz on portfolio optimization.

Listen to this article

Speed + | -



00:00 / 28:24

Feedback Listen on [Noa](#)

In current practice, however, executives in large corporations are reluctant to propose and advocate for risky projects. They quash new ideas in favor of marginal improvements, cost-cutting, and “safe” investments. Research studies long ago established this pattern. In a classic HBR article, for example, Syracuse University professor Ralph O. Swalm presented the results of a remarkable study of risk attitudes among 100 executives. He concluded that the findings “do not portray the risk-takers we hear so much of in industrial folklore. They portray decision-makers quite unwilling to take what, for the company, would seem to be rather attractive risks.” Our research confirms that this pattern persists.

In this article, we examine the phenomenon of risk aversion and avoidance and demonstrate how corporate incentives and decision-making practices exacerbate the problem. We present an analysis of just how much value executives leave on the table as a result and offer suggestions for mitigating the bias toward low-risk investments.

The Psychology of Loss Aversion

Economists and psychologists have long been aware that decision makers tend to place greater weight on the economic losses that could result from their decisions than on the potential equivalent gains. In 1979, Daniel Kahneman (a coauthor of this article) and the late Amos Tversky brought that concept to the forefront of management practice. (Their pioneering work in behavioral economics won Kahneman the 2002 Nobel Prize for economics.) Scores of empirical studies and experiments have further demonstrated the prevalence of loss aversion and identified its key features.

In a 2012 McKinsey global survey, for example, two of us (Koller and Lovallo) presented the following scenario to 1,500 managers: You are considering a \$100 million investment that has some chance of returning, in present value, \$400 million over three years. It also has some chance of losing the entire investment in the first year. What is the highest chance of loss you would tolerate and still proceed with the investment?

A risk-neutral manager would be willing to accept a 75% chance of loss and a 25% chance of gain; one-quarter of \$400 million is \$100 million, which is the initial investment, so a 25% chance of gain creates a risk-neutral value of zero. Most of the

surveyed managers, however, demonstrated extreme loss aversion. They were willing to accept only an 18% chance of loss, much lower than the risk-neutral answer of 75%. In fact, only 9% of them were willing to accept a 40% or greater chance of loss.



© Tim Booth/Stocksy United

What's more, the size of the investment made little difference to the degree of loss aversion. When the initial investment amount was lowered to \$10 million, with a possible gain of \$40 million, the managers were just as cautious: On average, they wouldn't make the investment if the chance of losing was higher than 19%. And once again, only 9% were willing to accept a chance of loss greater than 40%. This

indifference to the size of the investment seems perverse, because a relatively small investment is unlikely to present an existential threat to the enterprise and should, therefore, give managers scope to assume more risk.

Why are managers in large, hierarchical organizations so risk-averse? Swalm's tentative conclusion was that corporate incentives and control processes actively discourage managers from taking risks—a conclusion he felt was supported when managers he interviewed acknowledged that although their risk aversion was bad for their companies, it was good for their careers. We share his belief. CEOs are evaluated on their long-term performance, but managers at lower levels essentially bet their careers on every decision they make—even if outcomes are negligible to the corporation as a whole.

Consider how most investment decisions are made. A team with an idea for an investment puts together a business case for the project, which is then presented to a capital projects committee made up of the top managers of the unit. The champions of the project explain how it aligns with the company's strategy and provide financial models that assess the shareholder value it will create. The committee makes a decision on the basis of whether it judges the financial models and their underlying assumptions to be plausible.

Now consider the fact that this committee probably evaluates relatively few investment proposals. It is not unlikely, therefore, that if it were to allow a greater probability of failure for its investments, few or none of its decisions in a given time period would end in a successful outcome. The managers making the decisions would be held accountable for those outcomes and their reputations—possibly even their jobs—would be at risk. For all but the largest investments, the consequences of project failure would be far higher for the managers than for the company as a whole.

If this is true, we would expect that senior executives will be more open to investing in small projects than lower-level managers are. And that does appear to be the case. In a recent workshop, Nobel laureate Richard Thaler asked 22 heads of magazines owned by a large publisher if they would accept a hypothetical 50-50 investment that would pay \$2 million to the parent company if it was successful or lose \$1 million if unsuccessful.

Only three said they would accept the investment; the rest declined. In contrast, the CEO “wanted them all” to be accepted; he had a broader view of the possibilities and risks and realized that when the investments were pooled together, the risk profile was much more attractive.

The Value Left on the Table

In economic theory, unless a failed investment would trigger financial distress or bankruptcy, companies should aim to be risk-neutral, because investors can diversify risk across companies. Pure risk neutrality is unrealistic, of course, even for CEOs. Like the rest of us, they don’t want to lose their job over one bad, very large investment. But for investments that don’t threaten the firm’s viability, CEOs tend to be (as Thaler found) relatively risk-neutral, not only because they consider the size of the investments relative to the company’s resources but also because they recognize that the overall risk of a diversified portfolio is lower than the average risk of individual projects.

Unfortunately, as we’ve shown, companies regularly forgo smart investments because of managers’ aversion to risk. Suppose that each of your company’s 20 product lines has an opportunity to invest \$10 million with a 50% chance of receiving \$30 million and a 50% chance of losing the full \$10 million. In other words, each investment has an expected value of \$5 million: $(\text{gross gain of } \$30 \text{ million} \times 50\%) + (\text{gross loss of } \$0 \text{ million} \times 50\%) - \text{initial investment of } \10 million . Under the typical investment process, each unit head is likely to pass up the opportunity despite the positive expected value because they aren’t willing to bear a 50% risk and the pain of losing \$10 million.

From the company’s perspective, that’s a profoundly dysfunctional outcome. If the risk types of all the investments are uncorrelated, the simple math of probability (applying standard probability tables) will quickly tell you that there is only a 6% chance that the company as a whole would lose any money at all. Additionally, there is a 41% chance of earning more than \$100 million (after deducting the \$200 million investment) and a 75% chance of earning at least \$40 million.

This is not just theory. A technology company we advised carried out an aggregation of all its projects and their risks. First, using standard deviation of expected returns, executives estimated the expected value of each project proposal and the risks associated with each. They then built portfolios of projects and identified the project combinations that would deliver the best balance of risk and return. Executives could see that portfolios in general had higher returns than most projects deemed “low risk” and much lower risk than most of the projects with the same higher return as the overall portfolio. Taking a portfolio approach allows you to accept high-risk/high-return projects that you might otherwise turn down and reject low-risk/low-return projects that you might otherwise accept.

By separating decisions from execution, you can tailor incentives appropriately.

So how much money is left on the table owing to risk aversion in managers? Let’s assume that the right level of risk for a company is the CEO’s risk preference. The difference in value between the choices the CEO would favor and those that managers actually make is a hidden tax on the company; we call it the risk aversion tax, or RAT. Companies can easily estimate their RAT by conducting a survey, like Thaler’s, of the risk tolerance of the CEO and of managers at various levels and units.

For one high-performing company we worked with, we assessed all investments made in a given year and calculated that its RAT was 32%. Let that sink in for a moment. This company could have improved its performance by nearly a third simply by eliminating its own, self-imposed RAT. It did not need to develop exciting new opportunities, sell a division, or shake up management; it needed only to make investment decisions in accordance with the CEO’s risk tolerance rather than that of junior managers.

Creating an Aggregated Investment System

How do we change the practices and incentives around investment decisions so that managers become less risk-averse? To put it more bluntly, how do we ensure that managers don’t make decisions on the basis of personal (or local) consequences should their investments fail?

Make risky decisions in batches.

The first step is to establish a process in which projects are evaluated simultaneously with others on the basis of their collective value and risk. Ideally, a company would apply a portfolio optimization model that incorporates risk correlations across potential investment projects, as did the tech company we cited above. This approach would identify the least-risky portfolio for an overall target rate of return and risk given the investment opportunities available.

A simpler approach is to rank all projects across the company on the basis of their expected net present value (NPV) or some version of it, such as PV/I (present value divided by investment). PV/I is a common return measure that will be familiar to most managers, regardless of the business unit they belong to. Here's how the approach might work. Let's assume a company has five business units, each with 10 projects needing investment, for a total of 50 projects. Each unit proposes its 10 projects, presenting a careful risk assessment and a range of possible outcomes.

The corporate staff then ranks the projects across the company from highest to lowest in terms of expected value, ignoring risk. They accept projects, starting with the most value-creating project and continuing down the list, adding up the investment amounts required. Once the maximum amount of spending the company is comfortable with is reached, all projects left on the list are turned down, regardless of the business unit they belong to.

Next, the corporate staff examines the overall risk profile of the accepted projects. If the risk types of the projects are largely uncorrelated, the overall risk of the portfolio will be lower than the risks associated with almost all the individual projects. If the risk of certain projects is correlated, increasing the overall riskiness of the portfolio, corporate staff can swap in less-correlated projects from the remaining options on the list.

This selection process may well result in an uneven allocation of investments. One business unit might have eight projects approved, while another might have only two greenlighted. That information is useful in its own right: If one unit regularly finds

itself without projects on the list, that could indicate that it might be better off as part of another company or that its strategy should be narrowed to focus on generating cash rather than pursuing growth through new projects.

Ranking should be done annually at the very least, and preferably more frequently, depending on the length of projects. One company we know makes most of its investment decisions during designated weeks throughout the year so that it is regularly evaluating portfolios. If projects need decisions outside the normal cycle, the corporate staff can show the impact on the overall portfolio of adding them.

It might be argued that ranking or annual optimization imposes a certain amount of rigidity on the organization or prevents managers from reacting quickly to new opportunities or information. That may be true in certain instances, but many companies have found work-arounds. Some set up reserves to fund unexpected initiatives. Others require stage-gating: If circumstances change or projects don't meet predetermined milestones, the funds allocated to them during the annual process are shifted elsewhere.

Corporate incentives and processes actively discourage managers from taking risks.

Large, multibillion-dollar companies with many more than 50 projects across nonhomogenous units can easily modify the approach to handle the added complexity. Suppose a company has 25 business units and most of the projects are relatively small. It might allocate resources to business units rather than projects. Each unit would submit several investment-opportunity tranches, each reflecting a different investment goal. For example, a unit might submit a request for a tranche of \$200 million just to "keep the lights on," a second tranche of \$150 million to maintain market share and growth, and a third tranche might provide \$100 million for new products or services or for enhancements to customer service. Each tranche would have an estimated value and risk profile. They would be ranked across the company, and some units would receive all three tranches, others two or one or none.

Companies could also adopt a hybrid approach that combines allocations to business units and critical strategic projects (particularly new projects that address potentially large threats or opportunities, for which some degree of risk aversion might make sense). Strategic projects, whether or not they belonged to a particular unit, would be included along with the ranking of the business unit tranches. This approach ensures that critical projects get the attention of corporate leadership and that their funding is considered in a corporate context. One way to distinguish between normal and strategic projects is to have the CFO, in concert with the CEO, determine a project size below which risk neutrality is the goal. Projects larger than the designated size would be considered strategic.


Of course, simply introducing batch processes isn't enough to fully counteract loss aversion. For corporate staff and executives to make good decisions, they need high-quality input from the units on the prospects and risks of the investment opportunities. A common understanding of risk types is especially important, as executives will look to minimize correlation of risks between individual projects in a portfolio.

Bring risk out into the open.

In our experience, few project teams perform explicit risk assessments. They usually present a project to management with a set of cash-flow projections. They might include upside and downside cases, but nothing too drastic will be shown, and the returns will be close to the base case. The idea is to sell the project to management, and too much discussion of risk could frighten the horses. At one company we advised, the corporate team had performed a sophisticated Monte Carlo risk assessment on an array of projects, creating an appearance that it was transparent about discussing risk. Yet when we looked closely at the Monte Carlo output, we realized that in every case, the reported probability of negative NPV was zero.

We recommend that companies consider four or even five scenarios to achieve a good understanding of the risks. They should also abandon the practice of presenting a base case and up-or-down options, as the base case can too easily be seen as the default or most likely option, resulting in insufficient consideration given to the other scenarios. An even number of scenarios is helpful, because it lessens the chance that the middle case will be viewed as the default.

The first step in the risk assessment is to estimate the overall probabilities of each outcome. Executives are often reluctant to do this, because assigning probabilities can appear imprecise or subjective—but subjective probabilities are better than none. And they will get better at it with experience. It is often useful to have a number of executives assign probabilities, particularly those not advocating for the project. They will have less at stake, may be more objective, and may have a broader set of experiences. When many executives assign probabilities, the range of outcomes tends to be more extreme, which can help trigger useful discussions. And of course, the project champion should not be responsible for deciding which probability is accurate—that is a recipe for disaster.



© Tim Booth/Stocksy United

In one large company we know with long time horizons, a decision science team developed forecasts for the business unit heads, whose compensation was based in part on the amount of capital they invested in projects. Guess what? The team invariably

was “guided” (by the unit heads, who had ultimate authority over the forecasts) to revise their projections upward, which meant that more projects were approved. This phenomenon is not uncommon.

Next, teams should explicitly identify the critical risk factors that influence outcomes. If a team was investing in a pioneering process plant, for instance, it would need to consider product price risk, environmental risk, technological risk, currency risk, and, of course, execution risk. This last is worth calling out, because most companies don’t explicitly factor in execution risk—that is, human error on the part of managers carrying out the project, such as slow decision-making that leads to missed deadlines. To be sure, execution is controllable, and individuals can be held accountable for it, which may be why many companies don’t explicitly consider it a risk factor. But no organization is free of human error, so it is important to factor it in. Not doing so at the outset makes it likely that after a failure, more of the blame than is warranted will fall on execution.

When risks are specified in advance and agreed to by the whole team, executives are better able to identify the causes of project failure (and success). They can more easily determine whether an investment decision was good or bad, regardless of the outcome, which in turn makes it easier to take risks in the first place. It is important not to penalize poor outcomes, only bad decisions. Confusing the two is a great part of what makes managers risk-averse, which brings us to our next point.

Make risk less personal.

The final step in lowering risk aversion is to reduce employees’ personal risk in proposing projects that are outside the box. The simplest way to do that is to reward people whose projects are approved by senior management, regardless of the ultimate outcome of the project. A more sophisticated and preferable approach is to decouple the decision to pursue the project from its execution.

In this approach, if a new plant fails to earn an adequate return because demand is lower than expected, the failure is attributed to the decision to build the plant. If the plant fails because the project leader made construction errors that led to higher costs, the failure is attributed to execution.

By separating the decision from the execution, you can assign accountabilities to different people and tailor incentives appropriately. Accountability for decisions can be attributed to senior executives or to members of an investment committee, who have an incentive to maximize the value of the portfolio without being overly concerned about the risk of a single project. The execution risks, such as the cost and time involved in getting a plant up and running, can be assigned to the project leader, whose risks are mostly under her control.

We find that it also helps to consider longer time frames when evaluating decision-making performance. Managers often have too few projects in a given year for any single one to be assessed accurately or even fairly. One investment bank we know pushes accountability for projects up the hierarchy so that senior executives are responsible for many projects in a single year. Those executives' bonuses are highly variable from year to year, depending on how their project portfolios perform. Managers' bonuses, by contrast, are based on the performance of the multiple teams they participate in and are stretched out over three years. The longer time frame allows failures to be offset by successes so that penalties for managers with poorly performing projects are less severe.

Recognizing the inevitability of—or celebrating—failures is another practice that enables a culture of risk-taking. W.L. Gore, for example, gives “Sharp Shooter” trophies to managers outside product development who kill projects by identifying potential snags that the project team overlooked. The project team then writes up what it learned from the experience and how it could have made the decision to kill the project faster.

Finally, smart companies always make postmortems an important element of the management system. One company we know conducted them on its acquisitions and found that while the strategies were often sound, the executives assigned to integrate the acquired companies and carry out the strategies often lacked the resources necessary to be successful. Postmortems can also prevent companies from penalizing executives who executed well even though the external environment didn't behave the way the company had hoped.

CONCLUSION

Managers can control how their own behavior shapes an investment decision. But outcomes depend greatly on other people's decisions—decisions by competitors, regulators, and consumers. They are also influenced by factors beyond human control: natural disasters, commodity price spikes, the economic cycle. That means there is a strong element of chance in any investment, for which it is unreasonable to hold managers accountable. At a certain point, therefore, companies need to switch from processes predicated on managing outcomes to those that encourage a rational calculation of the probabilities. It's a switch that will deliver quick returns: Organizations that make inconsistent risk choices up and down the corporate hierarchy are leaving a lot of free money on the table.

A version of this article appeared in the March–April 2020 issue of *Harvard Business Review*.

Dan Lovallo is a professor of business strategy at the University of Sydney Business School and a senior adviser to McKinsey & Company.

Tim Koller is a partner at McKinsey in Stamford, Connecticut.

Robert Uhlener (robert_uhlener@mckinsey.com) is a senior partner at McKinsey in San Francisco.



Daniel Kahneman is the Eugene Higgins Professor of Psychology, Emeritus, and Professor of Psychology and Public Affairs, Emeritus, at Princeton University. He was awarded the Nobel Prize in Economic Sciences in 2002 for his work (with Amos Tversky) on cognitive biases.

This article is about RISK MANAGEMENT

[⊕ Follow This Topic](#)

Comments

Leave a Comment

Post Comment

0 COMMENTS

▼ [Join The Conversation](#)

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.