



CORPORATE FINANCE PRACTICE

Overcoming a bias against risk

Risk-averse midlevel managers making routine investment decisions can shift an entire company's risk profile. An organization-wide stance toward risk can help.

**Tim Koller,
Dan Lovallo, and
Zane Williams**

Here's a quick test of your risk appetite. Your investment team has approached you with two variations of the same project: you can either invest \$20 million with an expected return of \$30 million over three years or you can invest \$40 million with an expected return of \$100 million over five years (and a bigger dip in earnings in the early years). In each case, the likelihood that the project will fail and yield nothing is the same. Which would you choose?

Much of the commentary about behavioral economics and its applications to managerial practice, including our own, warns against overconfidence—that biases in human behavior

might lead managers to overstate the likelihood of a project's success and minimize its downside.¹ Such biases were certainly much debated during the financial crisis.

Often overlooked are the countervailing behavioral forces—amplified by the way companies structure their reward systems—that lead managers to become risk averse or unwilling to tolerate uncertainty even when a project's potential earnings are far larger than its potential losses.² In fact, the scenario above is based on the experience of a senior executive in a global high-tech company who ultimately chose the smaller investment with the lower up-front cost. That variation of the

project would allow him to meet his earnings goals, and even though the amount of additional risk in the second variation was small—and more than offset by a fivefold increase in the net present value—it still outweighed the potential rewards *to him*.

For projects of this size at a large company, the profit forgone by choosing a safer alternative—putting less money at risk with a shorter time to payoff—is modest: in this case, about \$20 million. But the scenario becomes more worrying when you consider that dynamics like this play out many times per year across companies, where decisions are driven by the risk appetite of individual executives rather than of the company as a whole. In a single large company making hundreds of such decisions annually, the opportunity cost would be \$2 billion if this were to happen even 20 times a year over five years. Variations of this scenario, played out in companies across the world, would result in underinvestment that would ultimately hurt corporate performance, shareholder returns, and the economy as a whole.

Mitigating risk aversion requires that companies rethink activities associated with investment projects that cause or exacerbate the bias, from the processes they use to identify and evaluate projects to the structural incentives and rewards they use to compensate managers.

A widespread challenge

The right level of risk aversion depends on the size of the investment. CEOs making decisions about large, unique investments are typically more risk averse than overconfident—and they should be, since failure would cause financial distress for the company.

In contrast, midlevel executives making repeated decisions about the many smaller investments that a company might make during the course of a year—expanding a sales force at a consumer-goods company into a new geography, for example, or introducing a product-line extension at an electronics firm—should be risk neutral. That is, they should not overweight negative or positive outcomes relative to their actual likelihood of occurrence. Decisions about projects of this size don't carry the risk of causing financial distress—and aversion to risk at this level stifles growth and innovation. Risk aversion is also unnecessary because statistically, a large number of projects are extremely unlikely all to fail (unless they are highly correlated to the same risks). Yet many managers at this level—who make many such investments over a career—exhibit an unwarranted aversion to risk.

In fact, we frequently run across CEOs stymied by their company's struggle with risk; decisions that may be in the best interest of individual executives, minimizing the risk of failure, are actually harmful for their companies. As the CEO at a manufacturing company observed, his company's business unit-level leaders gravitate toward relatively safe, straightforward strategies with earnings goals that seem reachable, even if these strategies mean slower growth and lower investment along the way. We have also heard from many nonexecutive board members that their companies are not taking enough risks.

Their anecdotal observations are consistent with findings we reported last year that suggested executives are as risk averse about small investments as they are about large ones.³ When we tested how 1,500 executives from 90 countries reacted to different investment scenarios, we discovered

that they demonstrated extreme levels of risk aversion regardless of the size of the investment, even when the expected value of a proposed project was strongly positive. Specifically, when presented with a hypothetical investment scenario for which the expected net present value would be positive even at a risk of loss of 75 percent, most respondents were unwilling to accept it on those terms. Instead, they were only willing to accept a risk of loss from 1 to 20 percent—and responses varied little, even when the size of the investment was smaller by a factor of ten. This is almost shocking, as it suggests that the level of risk aversion is remarkably constant within organizations, when it should vary based on the size of the investment and its potential to cause financial distress.

Understanding the source of risk aversion

Much of the typical risk aversion related to smaller investments can be attributed to a combination of two well-documented behavioral biases. The first is *loss aversion*, a phenomenon in which people fear losses more than they value equivalent gains. The second is *narrow framing*, in which people weigh potential risks as if there were only a single potential outcome—akin to flipping a coin only once—instead of viewing them as part of a larger portfolio of outcomes—akin to flipping, say, 50 coins. Together, these two biases lead to a distinctive set of preferences outlined in Daniel Kahneman and Amos Tversky's prospect theory, which was largely the basis for Kahneman's 2002 Nobel Prize in Economics.⁴

Consider a simple example of a risk-averse manager⁵ weighing whether to invest \$50 million today in a project that has an equal likelihood of returning either \$100 million or \$0 a year from now. If we were to ignore the time value of

money, we would expect a risk-neutral manager to be indifferent to the project—because the potential gains are equal to the potential losses. If the upside were greater than \$100 million, we would expect the same manager to make the investment. However, the upside would have to be almost \$170 million to entice the typical risk-averse manager to make the investment. In other words, the upside would have to be about 70 percent larger in order for that manager to overcome his or her aversion to risk.

But what if we were to pool these risks across multiple projects? If the same manager faced not one decision but ten, the story would change. The manager's range of outcomes would no longer be an all-or-nothing matter of success or failure, but instead a matter of various combinations of outcomes—some more successful, some less. In this case, the same manager would be willing to invest if the upside were only \$103 million, or only 2 to 3 percent above the risk-neutral point. In other words, pooling risks leads to a striking reduction in risk aversion.

Many of the managerial tactics used by companies in their capital-allocation and evaluation processes fail to take note of these basic behaviors. By considering the success or failure of projects in isolation, for example, they fail to understand how each will add risk to the company's overall portfolio and institutionalize a tendency toward risk aversion, essentially recreating the narrow framing that occurs at the individual level. To make matters worse, many companies also hold individuals responsible for the outcomes of single projects that have substantial uncertainty and fail to distinguish between "controllable" and "uncontrollable" events, leaving people accountable for outcomes they cannot

influence. As a result, many companies wind up with risk aversion at the corporate level that resembles that at the individual level—squandering the risk-bearing advantages of size and risk pooling that should be one of their greatest strategic advantages. In fact, many companies seem to exacerbate loss aversion, which is the primary driver of risk aversion.

Toward a company-wide approach to risk

Companies can reduce the effects of risk aversion, where appropriate, by promoting an organization-wide attitude toward risk that guides individual executive decisions. More specifically, companies should explore the following:

Up the ante on risky projects. Risk-averse organizations often discard attractive projects before anyone formally proposes them. To encourage managers and senior executives to explore innovative ideas beyond their comfort levels, senior executives might regularly ask them for project ideas that are risky but have high potential returns. They could then encourage further work on these ideas before formally reviewing them. They could also require managers to submit each investment recommendation with a riskier version of the same project with more upside or an alternative one.

Consider both the upside and downside.

Executives should require that project plans include a range of scenarios or outcomes that include both failure and dramatic success. Doing so will enable project evaluators to better understand their potential value and their sources of risk.

These scenarios should not simply be the baseline scenario plus or minus an arbitrary percentage. Instead, they should be linked to real business

drivers such as penetration rates, prices, and production costs. For example, when evaluating the introduction of a new consumer-goods product, managers should explicitly consider what a “home run” scenario would look like—one with high market share or high realized unit prices. They should also look at a scenario or two that captures the typical experience of product introductions, as well as one scenario where it flops. By forcing this analysis, executives can ensure that the likelihood of a home run is factored into the analysis when the project is evaluated—and they are better able to thoughtfully reshape projects to capture the upside and avoid the downside.

Avoid overcompensating for risk. Managers should also pay attention to the discount rates they use to evaluate projects. We repeatedly encounter planners who errantly use a higher discount rate simply because an outcome is more uncertain or the range of possible outcomes is wider (see “Avoiding a risk premium that unnecessarily kills your project,” on mckinseyquarterly.com). Higher discount rates for relatively small but frequent investments, even if they are individually riskier, do not make sense once projects are pooled at a company level.

Instead, if companies are concerned about risk exposure, they might adopt a rule that any investment amounting to less than 5 to 10 percent of the company’s total investment budget must be made in a risk-neutral manner—with no adjustment to the discount rate.

Evaluate performance based on portfolios of outcomes, not single projects. Wherever possible, managers should be evaluated based on the performance of a portfolio of outcomes, not punished for pursuing more risky individual

projects. In oil and gas exploration, for example, executive rewards are not based on the performance of individual wells but rather on a fairly large number of them—as many as 20, in one company. Hence, it may not be surprising to find that oil and gas executives pool risks and are more risk neutral.

Reward skill, not luck. Companies need to better understand whether the causes of particular successes and failures were controllable or uncontrollable and eliminate the role of luck, good or bad, in structuring rewards for project managers. They should be willing to reward those who execute projects well, even if they fail due to anticipated factors outside their control, and also to discipline those who manage projects poorly, even if they succeed due to luck. Although not always easy to do, such an approach is worth the effort.



The corporate center must play an active role in implementing such changes—in setting policy, facilitating risk taking, and serving as a resource to help pool project outcomes. It will need to become an enabler of risk taking, a philosophy quite different from that currently expressed by many corporate centers. The office of the CFO should also be involved in oversight, since it is particularly well suited to serve as manager of a company’s portfolio of risks, making trade-offs between them and taking a broader view of projects and the effects of risk pooling. ○

- ¹ Daniel Kahneman and Dan Lovallo, “Delusions of success: How optimism undermines executives’ decisions,” *Harvard Business Review*, July 2003, Volume 81, Number 7, pp. 56–63.
- ² Daniel Kahneman and Dan Lovallo, “Timid choices and bold forecasts: A cognitive perspective on risk taking,” *Management Science*, 1993, Volume 39, Number 1, pp. 17–31.
- ³ Tim Koller, Dan Lovallo, and Zane Williams, “A bias against investment?,” *mckinseyquarterly.com*, September 2011.
- ⁴ Daniel Kahneman and Amos Tversky, “Prospect theory: An analysis of decision under risk,” *Econometrica*, 1979, Volume 47, Number 2, pp. 263–91.
- ⁵ That is, a manager with a standard concave utility curve of the type $U(x) = x^{.575}$ in the domain of gains.