

Fretting about Modest Risks Is a Mistake

Matthew Rabin^{1,2} and Max Bazerman²

California Management Review
2019, Vol. 61(3) 34–48
© The Regents of the
University of California 2019
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0008125619845876
journals.sagepub.com/home/cm



SUMMARY

Managers often engage in risk-averse behavior, and economists, decision analysts, and managers treat risk aversion as a preference. In many cases, acting in a risk-averse manner is a mistake, but managers can correct this mistake with greater reflection. This article provides guidance on how individuals and organizations can move toward greater reflection and a more profitable aggregate portfolio of decisions. Inconsistency in risk preferences across decisions is a costly mistake for both individuals and for organizations.

KEYWORDS: insurance, decision making, decision biases, risk management, risk

Imagine that you have been offered the chance to make money on a coin toss. Would you take the bet if you would win \$1,100 if it came up heads, but have to pay \$1,000 if it came up tails? Now suppose that you manage a \$2-billion business. Would you accept a new project that would give your firm a 50% chance of improving its profits by \$1.1 million, but give it a 50% chance of losing \$1 million? The risky choices yield an average gain of \$50 and \$50,000. Yet most people would avoid these risks, despite their attractiveness on an expected-value basis.

We often care not only about what might happen on average, but also what would happen if chance favors or disfavors us. Risking a 10% chance of paying \$100 to repair an item does not feel the same as paying \$10 for an extended warranty that eliminates the risk. Similarly, a project in your business that has a 20% chance of making \$5 million does not feel the same as one that guarantees \$1 million. Everyone (we assume) prefers more money, on average, but how we feel about risk tends to vary by circumstance and by person. That is, we have different “risk attitudes.” *Risk neutrality* is when you seek to maximize expected value, or what you will get on average. *Risk aversion* refers to the tendency to accept a lower expected value in return for reducing your risk. *Risk-seeking* behavior involves accepting a lower expected value by *adding* risk. When you pass on a 50/50 bet of winning \$1,100 or losing \$1,000, you are being risk averse; when

¹Economics Department, Harvard University, Cambridge, MA, USA

²Harvard Business School, Boston, MA, USA

you choose a 50/50 bet of losing \$2,100 or nothing over losing \$1,000, you are acting in a risk-seeking manner. Notice that risk-averse and risk-seeking behavior is not a matter of rejecting or accepting risk; instead, it is about rejecting or accepting risk in a way that conflicts with expected value. So, when you invest \$1,000 with a 20% chance of obtaining \$10,000 and an 80% chance of winning nothing, this is not risk seeking, because it has a positive expected value. A risk-neutral person would also take that bet.

Economists have traditionally argued that people are entitled to their own risk preferences, believing that how people balance their desire for more money on average with uncertainty depends on individual tastes and circumstances. On this matter, economists have tended to agree with the rest of us: people are allowed any intrinsic feelings they want. In fact, many strands of research in recent years, most prominently Kahneman and Tversky's prospect theory, have highlighted when risk-loving, risk-neutral, and (by far the most common) risk-averse behavior are most likely to occur.¹ But closely aligned with this line of enquiry, researchers have also shown that the intuitive, straightforward, one-at-a-time way people react to the risk leads them astray. The type of framing effects and context dependence intrinsic to prospect theory and other empirical regularities mean that people don't realize that what they are doing in sum is likely not what they want in the aggregate. And Rabin formalizes the accumulated realization by economists that these descriptions cannot hold up to scrutiny as good advice when thinking about all the small to moderate risks that individuals avoid over time.² We follow Koller, Lovallo, and Williams, who apply similar concepts to focus on aggregation of managerial decisions.³ Our contribution focuses on highlighting how excessive risk aversion is built into individual cognition, that the problem aggregates in organizations, that risk inconsistency is a problem independent of risk aversion, and that behavioral evidence provides us insights into the cognitive source of these problems. We also highlight how our recommendations are at odds with human intuition and typical managerial practices.

The vast majority of times when people pay attention to risk, and not solely expected value, they are making a mistake. Even without knowing you or your business, we are fairly certain that you often pay to avoid risk (and at times pay to *add* risk). We believe you should stop buying extended warranties and forms of insurance where you could afford a loss, and more broadly quit worrying about risk in lots of realms where the worst outcome wouldn't be terrible. Nor should you give away expected value and take on risk just to avoid a loss. Just do what's best on average, and you'll gain the benefits from your willingness to bear risks—that is, all those premiums that you didn't pay for warranties that you didn't need.

Given that we believe that people are wise to follow their own preferences, why do we say most people make mistaken risk choices? Our argument is not that you should change your true risk preferences. Instead, it is that intuitive decisions you tend to make based on these feelings are likely a mistake *in terms of your own goals*. When you deviate from risk neutrality, you are losing more value in the long run than you may realize—and, worse, not reducing your risk as much as you may

realize. It turns out that, from a mathematical and statistical perspective, given your own risk feelings, you ought to be nearly risk neutral when making most of life's choices that do not involve huge stakes. Although we are pretty sure this is true for individuals, it is complicated by the fact that you may want to accommodate more emotional reactions to risk. But the case is even stronger for organizations—it seems especially unlikely that such responses should be accommodated. As a manager, you should insist that your employees pursue risk neutrality—that is, do what's best on average, even if that means alarming the lawyers.

And we have a second related argument: You and your employees' risk attitudes ought to be more consistent across circumstances and across different units in your organization than intuitive responses likely would lead them to be. Risk consistency means having the same risk preferences across decisions, and risk inconsistency is very expensive. Indeed, as loath as economists and business professors are to question people's personal tastes and goals, most of these experts have an intuitive predilection toward risk neutrality for non-huge stakes, and toward risk consistency when you can't or shouldn't simply be risk neutral. In this article, we explain why this is so, and why we agree with their predilection.

The Case for Personal Risk Neutrality

Why is it wrong to fret about each risk we face? We'll try to answer that from several different angles. Most of what we say, you'll see, boils down to the logic of adding up the outcomes from individual risky decisions: if you make decisions with more probable and better good outcomes than the probability of big bad outcomes, then all the good outcomes taken together will likely cancel out all the bad outcomes taken together. So if you ignore risk when making individual choices that do not involve huge stakes, you will not only do better on average, but face very little downside risk over the long term.

The Power of Aggregation

To clarify the cumulative effect of fretting about risk on a case-by-case basis, let us start by examining what happens when you think of any particular decision as if you were making it over and over again. In a famous example, in 1963, Nobel Prize-winning economist Paul Samuelson reported offering a colleague in the MIT Economics Department a bet involving a flip of a coin: the colleague would either gain \$200 or lose \$100. The colleague declined the bet, but said he would accept 100 such bets, based on 100 independent coin flips altogether.

A bit of arithmetic illuminates the attractiveness of the aggregate bet. To start, suppose you knew you would win exactly 50 times and lost exactly 50 times. Then you will gain exactly \$5,000: you are getting \$10,000 = 50 x \$200 from the 50 times you win, and losing \$5,000 = 50 x \$100 from the 50 times you lose, so you are getting \$5,000 on net. So if you were guaranteed exactly 50 wins and 50 losses, you would obviously accept 100 bets and take the \$5,000.

But you have no such guarantee; the bet is risky. Indeed, almost half the time, you'll lose more than 50 of the flips. Here's where some statistics come in. Sure, you might lose more than 50 times. But you would need to lose at least two-thirds of the 100 flips to come out behind given the 2:1 payoff. And that is far less likely than you may think. There is only a 1 in 700 chance that you'll lose any money (by flipping 67 or more tails), and a 1 in 25,000 chance that you will lose more than \$1,000.

It is doubtful Samuelson's colleague did those calculations, but as an MIT professor, he probably had a pretty good intuition for how favorable the gamble was. The rest of us are not so good at making such calculations. Indeed, Benartzi and Thaler interviewed evening MBA students, coffee shop visitors, and undergraduates about how they felt about taking both one of the bets offered by Samuelson and 100 such bets based on independent coin flips, aggregated.⁴ The percentage of MBA students willing to take the bet increased from 64% for one of the bets to only 75% when repeated 100 times, coffee shop visitors increased from 43% to only 66%, and undergraduates *decreased* from 77% to 50%. Clearly, these people did not fully appreciate the statistics behind repeating the gambles.

To verify the role of bad statistical intuitions in people's unwillingness to take even the aggregate gambles, Benartzi and Thaler ran a more complicated real-stakes experiment with undergraduates. They were offered the opportunity to play 150 independent gambles that would be favorable on average, but where each gamble risked a loss. Although the aggregate bet would yield them \$6, on average, with only a 1/300 chance of losing any money, only 49% of those offered the gamble accepted it. Benartzi and Thaler showed that the failure to understand the statistics was the prime culprit in this low take-up rate. They found that, on average, participants believed their odds of losing money across the 150 gambles was 1/4 rather than the true 1/300. Indeed, when participants were shown in pictorial form just how favorable the gamble was, the number willing to take the aggregate gamble jumped to 90%. The key conclusion is that risk intuition is faulty, and highlighting the pattern of common errors can lead to wise intervention.

People often may, like these smart undergraduates, reject a large number of risky (independent) options because they don't realize just how un-risky a proposition it is. As considerable research shows, very few of us fully appreciate "The Law of Large Numbers," which shows that through repetition of chance events, we are very unlikely to end up far from average. This holds so long as the events are reasonably "independent," meaning that getting good or bad outcomes on one does not affect the odds of good or bad outcomes on others. As Benartzi and Thaler and others have shown, a big consequence of our failure to appreciate the Law of Large Numbers is that we worry too much about the financial risk of repeated choices.

But what's wrong with being risk averse if you are only betting on *one* flip of the coin? And what's the problem with buying *one* extended warranty? Well, across your life span, you will make thousands of risky financial decisions.

Samuelson's colleague, savvy enough to see the statistical logic that most of us don't notice, still didn't manage to treat his isolated choice for what it was: part of the massive collection of risky choices he faced in his life. True, these risks usually won't be identical. The math is more complicated when each risk is different, but the basic relevance of large numbers is the same. Aggregated across all of those small-to-medium decisions, you are very likely to be better off if you simply take all positive-expected-value risks and reject all negative-expected-value bets. If you recognize that you face a huge number of small risks in life, rather than just the one staring you in the face, then your decision should more or less always be based on the aggregated version of the gamble. And if that's how you think about it (and you understand the Law of Large Numbers!), you are likely to be more risk neutral.

Inconsistent Risk Preferences

Most of us tend to act risk averse in some contexts and risk seeking in others. Because we do so in ways that follow our risk feelings in these different contexts, most people see no problem with that. We want to change your view. Consider the following adaptation of a classic set of problems that psychologists Amos Tversky and Daniel Kahneman posed to survey participants in 1986.⁵

Imagine that you face the following pair of concurrent decisions. First examine both decisions, then indicate the options you prefer.

Choose between:

- a. A sure gain of \$240, or
- b. A 25% chance of gaining \$1,000, and a 75% chance of gaining nothing.

Choose between:

- c. A sure loss of \$750, or
- d. A 75% chance of losing \$1000, and a 25% chance of losing nothing.

Prospect theory predicts people will pick (a) over (b), since people tend to be risk averse regarding gains, and to pick (d) over (c), since people tend to be risk seeking regarding losses. As Tversky and Kahneman predicted, 84% of their participants chose (a) over (b), and 87% chose (d) over (c).

Here's the problem with these choices: The combination of (a) and (d), chosen by most participants, adds up to a 75% chance of losing \$760 and a 25% chance of gaining \$240. But the combination of (b) and (c), which only 3% of people chose, leads to a 75% chance of losing \$750 and a 25% chance of gaining \$250. The combined option that a strong majority of people chose is clearly worse than the combined options that almost no one chose. Tversky and Kahneman assumed that their participants did not actually prefer (ad) to (bc); instead, the participants simply did not think about the combined choice. Knowing that it was obvious that no one would ever *intentionally* choose (ad) over (bc), Tversky and

Kahneman did not ask people directly. But in the process of replicating these results (using far smaller stakes), Rabin and Weizsäcker *did* ask people this question.⁶ In seeming violation of one of the first principles that every experimental psychologist and economist knows—that on any particular question at least 10% of participants will give some goofy answers—they found that exactly zero chose (ad). The choices are clearly because participants don't see the aggregate consequences of isolated risks. This happens despite the fact that Tversky and Kahneman quite dramatically put a pair of choices on the same sheet of paper! In life, choices are rarely so easy to compare.

As Rabin and Weizsäcker show mathematically, given the multitude of choices we make in life, the one-choice-at-a-time approach to choices inevitably lead to a dramatic consequence: unless we approach every new decision with essentially the same risk attitude, there is almost surely a way to change some of our choices to get more money while not increasing our risk at all. In Tversky and Kahneman's problem, where people go from risk averse in one question to risk seeking in the other, the inconsistency is stark, but any time our risk attitudes vary, we are behaving in the same way. Whereas our earlier examples show that being consistently risk averse regarding small choices throws away a lot of money while barely lowering risk, these examples show that inconsistent risk preferences lead us to throw away money while not lowering our risk at all! As you'll see later in this article, organizations of diverse people with diverse incentives are likely even more prone to such inconsistency than individuals, and hence more prone to throwing away money needlessly.

Ignoring Probabilities

When considering any choice, you should of course pay attention to the possibility of a very bad outcome. Unfortunately, people tend to ignore the *probabilities* of bad outcomes. Although ignoring likelihoods often arises organically from our fears, this is an especially big problem because of the many people out there making money from these fears by distracting us away from how low the probabilities are.

These errors are most obvious in an area of the economy where excessive risk aversion is common: Insurance. Given that insurance companies make a profit, on average, after paying business operation costs, some fraudulent claims, and the cost of claimant lawyers, we can safely say that those who buy insurance are making bets with negative expected value. On top of that, insurance companies also know the actuarial statistics better than you do. Hence, virtually all purchases of insurance are acts of risk aversion—people should know they are giving up expected value. Although people often buy insurance not realizing the bad actuarial value, the degree of risk aversion that we often exhibit when buying insurance can be stunning. This can be for some of the reasons above, but it often occurs because we pay too little attention to the odds.

If you focus on worst-case scenarios—which we tend to do when an insurance salesperson is trying to sell us something to avoid those scenarios—you will

often fail to pay sufficient attention to how unlikely the terrible event is. When one of us was connecting his landline telephone after a move many years ago, a representative from the phone company tried to sell him internal wiring protection—a type of extended-service insurance of dubious actuarial value that, as it happened, the customer was then writing about. “Are you sure it’s a good idea? What’s the chance that I’ll need the services you’re asking me to pay for?” he asked. She replied, “Well, I don’t know how often it happens. But I can tell you that those who end up needing repairs to their internal wiring sure are glad they had the insurance!” If we can be persuaded to buy insurance based solely on whether it is good to have the insurance *if* the worst-case scenario happens and ignore the odds it will happen, we are going to buy too much insurance.

The art of getting you to ignore how unlikely bad outcomes are is especially impressive for those selling you extended warranties—it is generally right after salespersons convince you to buy a product by telling you about its remarkable reliability that they then tell you vivid tales of expensive repairs that you need to insure against. Yet extended-service warranties are one of many insurance products that tend to have amazingly low loss-ratios. In insurance language, a loss ratio is the percentage of the premiums collected that actually go to paying out claims. If you were risk neutral, you would only want a policy with a loss ratio greater than 100%, but of course, no insurance company would offer such a policy. Extended warranties have an approximate loss ratio of 20%, and travel and car-rental insurance have loss ratios closer to 40% and 10%, respectively. A key way people sell insurance products with very low loss-ratios is to persuade you that such a loss is more likely than reality dictates.

It is true that sellers of these products often prey on consumers who are not financially savvy. But we predict that even many readers of the *California Management Review* make far more risk-averse insurance decisions than makes sense. Do you carry comprehensive insurance on your car (to cover theft and damage to your car)? Most wealthy drivers do, even those who can easily afford to run a small risk of needing to pay for car repairs. Do you pay a higher premium for a lower deductible on your homeowners insurance? If so, you’re paying an extra premium to insure small losses that you can certainly afford to pay in the unlikely event that such a loss occurs.

Of course, some insurance is wise for virtually all of us, and our logic about how small- to medium-sized risks tend to cancel each other out does not apply to the common practice of paying a bit of expected value to avoid life-changing losses. People carry liability insurance on cars, which has negative expected value, to ensure that we can lead the rest of our lives as planned even if we cause a car crash with massive, and massively expensive, injury to others.

While it is sensible to worry about life-changing risk, even in this realm we often fret too much. One of us recalls meeting with a lawyer to draw up his will, along with his spouse. The lawyer recommended a trust that would cost \$4,000 extra. He asked the lawyer what the trust would protect, and the best the lawyer could come up with was that it would protect one spouse in the event that that

spouse lost a multi-million-dollar lawsuit in the same year that the other spouse died. We decided to follow the expected-value-maximizing strategy and reject this lawyer-earnings-maximizing idea. As this story conveys, lawyers often tell us how to reduce risk without doing the appropriate cost-benefit analysis of the risk they portray, encouraging very risk-averse behavior that is a bad idea no matter what the stakes are.

Creating Risk When You Think You Are Avoiding It

Some readers might think that such products have a market because less wealthy consumers cannot afford losses of \$1,000, and therefore need to insure against such relatively small losses and buy extended warranties. Yet these policies may not even reduce overall risk for such consumers. Buying an extended warranty on your car provides some protection against needing to pay a \$3,000 repair out of pocket. But if a consumer regularly buys insurance products with low loss-ratios because she cannot afford the losses, she may well have less funds available to cover other surprises that come up in life (such as eviction notices, tuition bills, and doctor's bills). These low-loss-ratio insurance products make certain risks in life salient, but as much as it seems that insuring specific losses is *avoiding* risk, in many cases, it is actually *adding* to risk. Savings offers the first and cheapest buffer against risk. When we reduce our savings by paying too much for expensive insurance, we not only lose pocket money, but *increase* our risk of being exposed to financial catastrophe. As Erica Jong once wrote (perhaps with a different context in mind), "And the trouble is, if you don't risk anything, you risk even more." Buying expensive insurance often increases the number of contingencies that you won't have the resources to handle.

The Emotions of Risk Taking

As confident as we are about the sheer statistical certainty that paying attention to risk for small stakes is going to lose you money in the long run, we should not be too glib in telling you how to live your life when confronting small risks. People might sensibly pay costs to avoid risk so as to ward off the adverse emotions that they anticipate a particular risk would make them feel. Imagine that you understood the loss ratio involved in car-rental insurance, and you could afford to pay for your rental car in the event of a loss. We have implied buying the insurance would be a mistake. But if you know that you would be miserable throughout your vacation if you drove without full insurance coverage, this may not be so clear. Köszegi and Rabin have argued that a real unhappiness in response to bad news can explain much of the risk avoidance we observe, without assuming mistakes.⁷

We might still urge a few tricks to feel the losses less painfully. Maybe try to think about what a given risk means solely in terms of the actual effect on your life—what will happen to your ability to buy things you want if you have more or less money—rather than the feelings triggered by the risk. Or, when possible, maybe try to avoid noticing the losses at all! Economist Richard Thaler often

advises checking the performance of riskier stock investments less often; that way, you suffer less on the 45% of days on which the stock market goes down.

Ultimately, however, if you are really bothered by the losses you experience, paying to avoid the risk could be worth the cost. What is far less likely, it seems to us, is that leaders of organizations should let their emotions guide them when doing so would lose money for the organization. So let us see how some of these lessons apply more directly to organizations.

Managerial Implementation

We have seen how you can be wiser by ignoring riskiness in favor of expected value as an individual. How can you make better decisions in your role as a leader? First, recognize that—with rare and recognizable exceptions (generally, risk that you cannot afford)—risk aversion is a mistake. Second, understand that risk preferences should be consistent across your organization. Thus, you should try to find a way to get everyone in your organization—including the risk seekers in sales and the risk avoiders in legal—to have a more consistent response to risk. Leadership involves providing guidance on thousands of decisions, many of them interconnected, across your organization. When these thousands of decisions reflect risk aversion—or, worse, inconsistency in dealing with risk—your organization is collectively giving away expected profit.

Finally, there are mind-set shifts that leaders should encourage in their organizations. There is no magic bullet available. Yet we believe that leaders need to consistently convey a series of arguments about what wise decisions look like, and then reward decisions consistent with this vision.

Risks Are Smaller for Shareholders

The argument for risk neutrality becomes even stronger when we consider the effect of organizational decisions on shareholders. Because shareholders typically hold stock in many firms, a risk that affects only one firm is not likely to pose much of a threat to their total portfolio. They can also buy lower risk investments to reduce their overall risk. A manager deciding whether to start a project with a 50% chance of increasing corporate profits by \$11 million and a 50% chance of decreasing profits by \$10 million should not think of it as a choice to get \$500,000 in expected value for the risk of losing \$10 million. The risky option does add up to a total of \$500,000 for shareholders, but no individual shareholder faces anywhere near \$10 million in risk. Leaders need to reward decisions that maximize expected value, even when the risky decision comes out negatively.

The Power of Aggregation within Organizations

One of the authors recently consulted for an insurance firm on its claims-settlement process. The firm spent many billions of dollars each year to settle claims, and it hired one of us to help improve the negotiation of settlements.

Claims adjusters within the company told the author stories in which his negotiation strategies frequently ran into the reality of the claims process. One common problem with multi-million-dollar claims concerned reinsurance—in this case, the practice of one insurer transferring portions of risk to other insurance companies by reinsuring the risk above a certain claim level. That is, Insurance Company A sells a policy to a client, and then turns around and buys an insurance policy from Insurance Company B to cover any settlement on the policy above \$5 million. Insurance Company B makes an expected value profit (and Insurance Company A makes an expected value loss) in return for Insurance Company B reducing Insurance Company A's risk.

When a claim occurs, negotiations begin over its value. Suppose the claimant is asking for \$20 million, and Insurance Company A has offered \$2 million. Imagine that the fair value of the claim (which is hard to assess) is \$6 million. A jury could award anywhere between \$1 million and \$20 million, with an expected award of \$6 million. Now let's imagine that Insurance Company A can settle the dispute for \$5 million. Will they? Perhaps not, since the worst they can do at trial is lose \$5 million, since Insurance Company B has to pay for the claim above this level. Obviously, Insurance Company A could negotiate with Insurance Company B to pay part of the \$5 million, since this would reduce risk for Insurance Company B (by settling, rather than going to trial), but this gets quite complicated. The point is that reinsurance can lead Company A to pass on a settlement that makes sense for the claim (for Insurance Companies A and B combined), because the asymmetric pattern of liability makes it in the interest of the lead insurance company (A) to not have an interest in settling.

This type of reinsurance through “tiers” (and sometimes multiple tiers) is common. When you ask why the reinsurance occurred in the first place, one sensible answer is that there might be highly correlated risks (e.g., a hurricane hitting an area and causing damage to multiple homes). But the most common answer is that it is meant to share risk. However, if the claims are reasonably independent, then we come back to the aggregation argument we made regarding individuals, just on a larger scale. Why would an insurance company that pays billions of dollars in claims want to reinsure with a company that seeks to make a profit for coverage in the \$5- to \$20-million range? From Insurance Company A's perspective, it is far wiser to keep the coverage for the wide variety of claims worth hundreds of thousands of dollars, millions of dollars, and even tens of millions. Reinsurance should be limited to truly vast or highly correlated claims that would put the future of the insurance company at risk. A company that pays billions of dollars in claims can earn higher profits by aggregating wise bets with positive expected value and avoiding the tendency to be risk averse in specific situations.

Because insurance companies make money based on the risk aversion of their customers, insurance is the starkest example we know of in which individuals and organizations should switch to more consistent, less risk-averse decisions. So, when we see insurance companies *themselves* forego expected profits by being risk averse, we become convinced that there is a fundamental problem in the way

firms deal with risk. Virtually all firms lose money to risk aversion. Indeed, many have entire departments that are institutionally risk averse—legal, credit, compliance, and so on. Like individuals, companies would be more profitable if they made the vast majority of their decisions with the goal of maximizing expected value, and they prevented individual employees and units from engaging in risk aversion (or risk seeking).

Passing on positive-expected-value-but-risky projects, hedging, buying too much insurance, and similarly overly cautious behaviors reduce organizations' expected value. If you imagine an organization as a set of thousands of risk decisions made by individual employees who may act in a risk-averse, risk-neutral, or risk-seeking manner, the leader of the organization should want to create an environment that promotes risk-neutral behavior for virtually all decisions. In the extreme, we could argue that the world "virtually" is not needed in the prior sentence, since the company's shareholders, for whom the stock probably represents a small portion of a diversified portfolio, should want the company to maximize the expected value of all its choices and investments. But, we accept the political realities of the owner-manager relationship, and can accept a small amount of risk aversion in cases where the existence of the firm is at stake.

A reminder is in order to prevent misapplication of the principle of aggregation. When talking about the 100 flips of a coin by Samuelson's colleague we stressed that these must be *independent* bets, and when we talked about the folly of reinsurance we noted that it is a sensible institution when an insurance firm is facing highly correlated risk (such as providing earthquake insurance to a large number of homes in the same place). One hundred bets on the same coin flip add up to a huge risk. And if all the risks throughout your organization hinge on the same event then you cannot rely on the Law of Large Numbers. Leaders need to remind their employees that they are making lots of decisions, and the organization is making far more decisions, and the goal is the optimal portfolio of decisions.

Inconsistency across the Organization

Perhaps the most striking problem we have identified, and the one that calls for the greatest intervention from leaders, is how inconsistency in risk responses can lead to throwing away money without risk reduction. Organizations tend to have inconsistent risk preferences across functions and product divisions. Salespeople often are implicitly risk seeking; they want to book the sale, even if a payment risk exists. Legal is typically risk averse, seeing its job as keeping the organization from accepting downside risk. We too often accept these varying risk attitudes as normal and acceptable. However, this state of affairs reduces the expected value of the organization. Sales should develop sales prospects and legal should provide legal advice, but all functions should maximize the organization's expected value—and this means correcting risk-seeking and risk-averse behaviors.

Evidence suggests that the type of consistency we recommend is far from a reality. Harvard Law School (HLS) students tend to focus heavily on avoiding

downside risk as opposed to maximizing average client profitability, former HLS student Michael Ghaffary found in his research. He also found that HLS students are far more risk averse than Harvard Business School students, and that risk avoidance increases throughout students' HLS experience. Of course, lawyers are responsible for advising clients about their legal risks. But the goal of that advice should be to help clients maximize their profitability in light of their legal risks. Risk aversion is a significant threat to this goal for all the reasons we have said, but even if you feel you must be risk averse, be aware that inconsistency in risk preferences across decisions in your organization could be very costly. And watch out for it in your own decisions.

Suppose you are a bigwig in your organization who has two back-to-back meetings this morning. In the first meeting, an employee tells you about an exciting new project that's sure to make big profit. You must decide between two variants of the project: a riskless, guaranteed \$70-million profit; or a slightly riskier version with a 50% chance earning \$64 million and a 50% chance of earning \$78 million. You get a huge profit either way, but the riskier one will give you an average profit of \$71 million, as compared to \$70 million. You might go for the risky option, and we would agree.

Now suppose that your second meeting of the morning is (so it seems) a riskier proposition. An employee asks you whether to pursue a \$6 million investment that might pay nothing or might return \$16 million. Thus, the investment offers a 50% chance of making \$10 million net and a 50% chance of losing \$6 million. Approving the project has an expected value of $(\$10 \text{ million} - \$6 \text{ million}) / 2 = \2 million . But because of the risk, you say no. That would make us frown (if you are a sizable firm that can afford such a loss), but we acknowledge that losing \$6 million is a lot.

If we couldn't persuade you to be risk neutral in both choices, however, as with one of the examples from above we would ask you to reverse your choice on the two decisions! Your morning's work for the firm is the sum of the two choices. Given your approval of the first risk and disapproval of the second, you have created a half-chance of making \$64 million and a half-chance \$78 million. Yet if you reversed your choices, you would have generated a half-chance of \$64 million (\$70 million minus \$6 million), and a half-chance of \$80 million (\$70 million plus \$10 million). With the reversal, you have the same downside risk, but a better upside. That sure looks like a better choice!

Both individual psychology and organizational factors tend to make our departures from risk neutrality quite varied: different decisions in the organization will variously be very risk averse, a bit risk averse, risk neutral, a bit risk seeking, or risk seeking for different decisions. When you add up the risks, the individual decisions become lost in the fold, and your organization only faces the total risk. In that light, you should strive to set a threshold for tolerating risk (again, ideally, we advise you to generally ignore the risk and focus on expected value) and apply it consistently. In many cases, it can be useful to bring together managers from different units—applying their different risk preferences—in order

to clarify the problem of losses that the organization accrues due to inconsistency across the organization.

Confronting Risk Aversion in Your Organization

Our goal is to help move your organization toward being risk neutral, risk consistent, and more profitable. To get there, we started with the easier of two tasks, as we have tried with our colleagues on many occasions: namely, convincing you to decide all of your small-to-medium-risk decisions on a risk-neutral, consistent basis. But we have not helped you much with the second, surely harder, task: managing risk wisely within your organization requires you to persuade hundreds or thousands of employees across the organization to implement decisions in a more risk-neutral, consistent manner.

To move in this direction, you will need to make sure that key professionals are trained in basic decision analysis and that they can see how deviation from risk neutrality makes the organization less effective. Leaders also need to make risk neutrality part of their organizations' culture. Employees must come to view deviations from risk neutrality as mistakes rather than as the expression of their different personalities or their department's function. When different divisions of the organization argue about risk, individuals should be reminded to make decisions with the long-term, risk-neutral preferences of the shareholder in mind.

It is also important to recognize that employees may be departing from risk neutrality because their incentives may be poorly aligned. Specifically, incentives sometimes induce very conservative behavior. Back when IBM was the dominant computer supplier, the phrase "No one was ever fired for buying an IBM computer" went around. At times, it likely would have been wise for an organization to experiment with buying cheaper alternatives to IBM computers. But if employees could expect to be punished if they bought a brand that failed, but not for overpaying for the IBM, then these incentives might have induced risk-averse behavior rather than maximizing expected value for the organization. At other times, incentives for competing with other employees could induce employees to be risk seeking rather than to maximize expected value.

Leaders would be wise to conduct a risk audit focused on better understanding risk preferences across the organization. Identifying the unique problems within your organization is relevant for tailoring solutions to your context.

A Few Caveats

We have little doubt most managers will stop any risk-seeking behavior when they can; whatever might be causing individuals within the firm to throw away expected value *and* increase risk, it is not good for the firm. But what do you do when you cannot stop all risk-seeking behavior? It might be tempting to balance risk seeking in one part of the organization with risk aversion in another, hoping this adds up to risk neutrality. It doesn't! It is more likely to add up to throwing away money without getting rid of the risk. Just as we pointed out earlier that being too risk averse can actually increase a person's exposure to risk by

lowering the buffer against other risks in the form of savings, so too a firm with some reckless elements who are taking unprofitable risks is often better advised to encourage prudent acts of value-maximizing risk neutrality over risk aversion, because such value maximization will create profits that can help shield the organization from those risks.

There are undoubtedly contingencies that affect the wisdom of the advice that we have offered, and contingencies that will affect the ease of implementing our recommendations. As we noted earlier, it may make complete sense to be risk averse when the life of the firm is truly at risk. And similarly, it is easy to develop a rational argument for individual risk aversion in contexts that cover true catastrophes. Similarly, when a group of medium-size risks are highly correlated, some risk aversion may be logical. Thus, we can imagine endorsing an insurance company buying reinsurance for covering massive hurricanes, where many policy holders may make moderate claims all at the same time for the same event.

It is not possible to eliminate risk aversion, but it is possible to reduce it dramatically in the many cases in which loss is affordable. Thinking through the actual implication of a loss, rather than the feelings triggered by the loss, can move leaders toward more rational risk preferences. Aggregation across uncorrelated risks, across the organization, will allow leaders to see the accrued savings that can be obtained.

Learn to Take the Hits for Long-Run Profits

Good decisions sometimes lead to bad outcomes, and bad decisions lead to good outcomes. When you look at a portfolio of outcomes over a long period of time, then that tells a pretty accurate story as to the wisdom of choices. But when you do not have the luxury of such evidence, managers need to look at the wisdom behind the decisions that are made, not their outcomes. Many companies pride themselves on being “results oriented.” However, this focus on results can lead to dysfunctional risk aversion if the manager can be punished for making a wise (e.g., maximizes expected value) decision that turns out poorly. For example, if the manager identifies a risky bet with a 20% chance of returning \$1,000,000 on a \$100,000 investment, there is an 80% chance that the decision will lead to a bad outcome. And, not all employees have the luxury of making a large portfolio of important decisions. A company should want all its employees taking these sorts of bets, and would be wise to reward those who make positive-expected-value bets, even when they turn out badly.

In many organizations, individuals’ risk choices are treated as personal preferences. Too often, these preferences are based on insufficiently considered practices. Individuals and organization should start their analysis without any risk aversion at all. They could then ask whether there is something extraordinarily unique about the specific decision to justify departure. This reorganization of when risk aversion is considered would lead to less risk aversion and a far better portfolio of decisions. Through greater risk neutrality, you and your organization can move toward greater profitability.

Author Biographies

Matthew Rabin is the Pershing Square Professor of Behavioral Economics in the Harvard Economics Department and Harvard Business School (email: matthewrabin@fas.harvard.edu).

Max Bazerman is the Jesse Isidor Straus Professor of Business Administration at Harvard Business School (email: mbazerman@hbs.edu).

Notes

1. D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47/2 (March 1979): 263-291.
2. M. Rabin, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68/5 (September 2000): 1281-1292.
3. T. Koller, D. Lovallo, and Z. Williams, *Overcoming a Bias against Risk* (New York, NY: McKinsey, 2012).
4. Shlomo Benartzi and Richard H. Thaler, "Risk Aversion or Myopia? Choices in Repeated Gambles and Retirement Investments," *Management Science*, 45/3 (March 1999): 364-381.
5. A. Tversky and D. Kahneman, "Rational Choice and the Framing of Decisions," *Journal of Business*, 59/4 (1986): 251-284.
6. M. Rabin and G. Weizsäcker, "Narrow Bracketing and Dominated Choices," *American Economic Review*, 99/4 (September 2009): 1508-1543.
7. B. Köszegi and M. Rabin, "Choices, Situations, and Happiness," *Journal of Public Economics*, 92/8-9 (August 2008): 1821-1832.