

SIEPR

policy brief

Stanford Institute for Economic Policy Research

on the web: <http://siepr.stanford.edu>

Bad Bets and Incentives for Risk-Taking

By *Orie Shelef*

As the run-up to the Great Recession demonstrated, risk taking by managers can have drastic consequences not just for their own firms but also for the global economy. Yet, risk taking is an integral and important part of firms and our economy. Sensible risk taking is central to economic growth, entrepreneurship, and innovation. The inherent tension is that while we do not like risk taking, there are projects that, on average, succeed by enough to justify it. We want to give people incentives to undertake such risky projects, but not those where the expected outcome does not justify the risk. This policy brief will not answer the fundamental tension. What this policy brief will do is complicate the problem by introducing another kind of risk taking that we can all agree we would be better off without—one, however, that seems to be

undertaken systematically in response to compensation schemes—namely, bad bets. In this brief, I discuss what are bad bets, why incentives can lead people to take them, empirical evidence that incentives actually do lead to bad bets, why this matters for all sorts of incentives, and what policymakers can do.

What are bad bets? And how do compensation plans encourage them?

Bad bets, in my terminology, are simply risky opportunities that provide more risk but not more expected return than some other available opportunity. Gambling is easily seen as a bad bet. On average, a gambler takes risk and loses money. If he just kept his money in his pocket, a gambler would have less risk and a higher expected return

continued on inside...

About The Author

Orie Shelef is a postdoctoral fellow at SIEPR through spring 2016. Prior to joining SIEPR, he attended the University of California, Berkeley, where he received a PhD. His applied microeconomic research focuses on organizations, including performance, risk taking, incentives, and structure in private and public settings, with a secondary research interest in energy and the developing.



SIEPR *policy brief*

(not losing anything). Gambling, as the house, however, is not obviously a bad bet. The house, on average, wins, but it does take some risk. Some bad bets might look pretty good—what makes bets bad is that they are dominated by a better alternative. For example, you can invest in a stock market index through a mutual fund that charges high fees or through one that charges low fees. Both may be expected to do better than other choices. But, assuming they make the same investments and provide you the same services, the high fee fund is a bad bet. You get more return on the low fee fund without taking on more risk.

Bad bets are different from the classic concerns about risk taking. Unlike the usual concern about finding the optimal trade-off between risk taking and performance, there is no trade-off. Choosing the wrong risk/return trade-off is inefficient, but in some sense second order, compared with bad bets. Also, bad bets are not normally part of the relevant choice set. Individuals do not like risk, and firms either do not like risk or are indifferent to it. They have no reason to even consider a riskier, worse-performing choice. Dominated, and thus irrelevant, choices should not matter. This intuition has driven the risk/

reward trade-off. More reward comes with more risk and more risk comes with more reward. As long as we are making decisions for ourselves, this is straightforward, but once we delegate the decisions—to executives, money managers, employees, etc.—those we delegate to have their own incentives. The classic research in this setting has grappled with whether the managers we use will pick the right risk/return trade-off for our risk tolerance—they might take too little risk because they have different risk tolerances than we do or have different incentives. Yet, as Palomino and Prat (2003) argue, we may pay them in a way that makes them like risk. If we do, then bad bets are relevant, and the whole logic of more risk/more return may not hold.

The central feature of compensation that can induce bad bets is when managers have more to gain from succeeding than they have to lose from failing. The simplest part of a compensation scheme that has this feature is a stock option. An option pays a share of the performance if the performance is above a threshold but does not penalize if performance is below a threshold. For example, consider a manager in a firm with a stock price of \$40 and an option—the right to buy

a share—at \$50. Suppose this manager is choosing between two projects, each with a 50 percent chance of succeeding: a good bet that makes the company worth \$49 if it succeeds and \$39 if it fails, an average return of 10 percent, and a bad bet that makes the company worth \$55 if it succeeds and \$17 if it fails, an average return of 10 percent. With either bet, the option is worthless in failure. Success with the good bet is also not rewarded by the option. But the option pays out \$5 if the bad bet succeeds. The manager's option encourages choosing this bad bet over a lower risk and higher expected return choice. But other explicit compensation, such as increasing commissions, or implicit rewards, such as those for exemplary performance, can have similar incentive effects as an option. And these types of compensation schemes are common. However, all of these features can also induce valuable risk taking, especially if individuals are more risk-averse than their firms. Thus, the importance of bad bets is not something we can infer just by looking at how people are paid. Instead, we have to look for evidence of people taking more risk but getting low returns.

Bad bets are being made, and they matter.

Researchers are just starting to find this evidence. Huang, Sialm, and Zhang (2011) and others look at mutual funds and puzzle that funds that start to take more risk do worse. Shue and Townsend (2013) look at executives that receive more options, encouraging more risk, and find that they take more risk. But they also note, without explanation, that their firms, at least in the short term, do worse. These both are consistent with taking bad bets but still leave open alternative explanations.

In my work (de Figueriedo, Rawley, and Shelef, 2013), we connect the two. We do that by exploiting, from a research perspective, hedge funds. Hedge funds provide a setting that we can use to disentangle the major challenge of studying incentives in the real world: Does a manager do A because of an incentive? Or did someone decide to offer that incentive because the manager was going to do A? Because of the nature of the hedge fund contracts, the fee formula is fixed in advance and we take advantage of market movements to change the incentives for risk in a way that was not chosen by the manager or by investors. And, we can address the major challenge of studying risk taking, by being

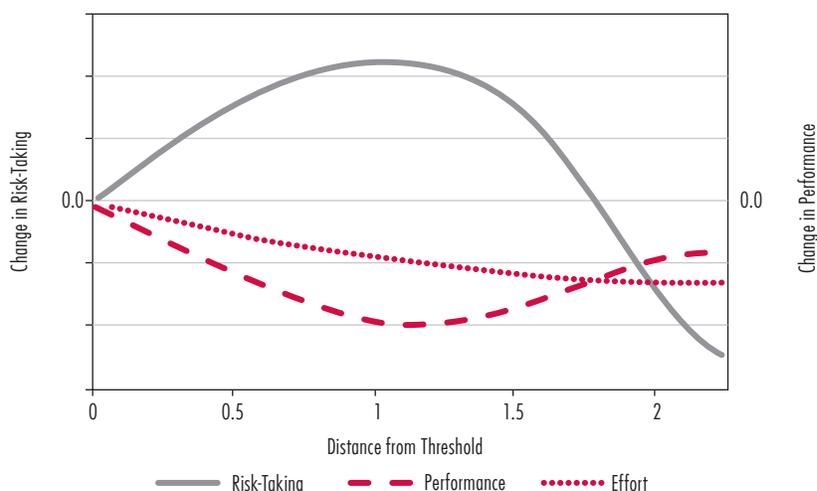
able to measure and compare risk taking. Since hedge fund managers are free to make all sorts of investments, they all could have made the same choices and looking at the risk embedded in their performance this lets us measure how much risk they ended up taking.

The key feature of hedge fund incentives that we exploit is that hedge funds have a threshold, called a “high-water mark,” which works like an option. If they are below this threshold, they get paid a management fee, say, 2 percent of the assets they manage. Above the threshold, they get the management fee plus they get a performance fee, usually 20 percent, of the returns they generate above the threshold. Thus managers

share in 2 percent of losses but 22 percent of gains. These are strong incentives to exert effort, but also to take risk. Yet, these incentives vary depending on where a manager is with respect to the threshold—for example, the \$50 option discussed above may affect incentives differently if the stock price were \$30, \$40, or \$60.

Figure 1 shows what we would expect to see, if managers are making bad bets. Each line compares the effort, risk taking, and performance of the manager, when the manager is some distance below the threshold compared with normal risk and performance levels. As managers move below the threshold, they exert less effort (dotted line). This is driven

Figure 1
Predicted Risk-taking, Performance, and Effort in Response to Incentives



by a bigger chance that the payoff they receive for this effort is 2 percent rather than 22 percent. At first, risk taking (solid line) increases. This is because taking more risk still pays off in the 22 percent region, but the cost of risk falls as it is more likely costing 2 percent rather than 22 percent. Eventually, as managers are far enough below the high-water mark, the risk stops paying off and managers take less risk, because they have little chance of reaching the 22 percent region but are still paying 2 percent of the cost of the bad bets.

Empirically, we cannot observe effort, but we do see performance (dashed line). At first both reduced effort and increased risk taking

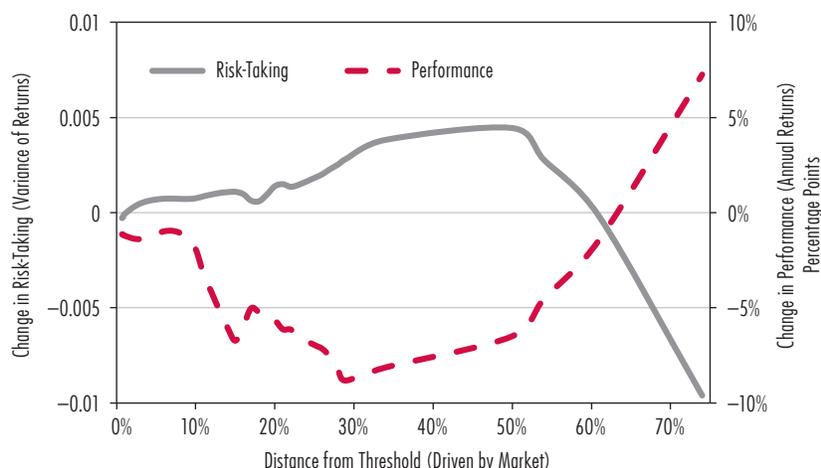
lead to reduced performance. Eventually, if bad bets are important to performance, then we should see managers very far from their thresholds stop taking bad bets and improve their performance.

Figure 2 shows what managers actually do, as market performance moves them away from their thresholds. The solid line plots risk taking and the dashed line plots performance. It looks much as Figure 1 predicts. At first, risk taking rises and performance drops, which is consistent with managers making bad bets, but it is also consistent with taking more risk and exerting less effort. When managers are very far from their thresholds, the effects of bad bets show up

clearly. Risk taking drops, as we would expect. Risk taking of any kind just does not pay much this far away. While we do not observe effort directly, managers far from the high-water mark should be exerting even less effort than those closer. Yet, their performance rebounds. The increase in performance relative to managers somewhat closer to the high-water mark is exactly what would happen, if those closer were taking bad bets and that was driving their performance decline.

Bad bets are economically significant and worthy of policy discussion. Extrapolating the pattern in Figure 2 to the entire industry, hedge fund investors lose \$20 billion a year and subject investors to risk equal to another \$12 billion.¹ But the underlying mechanism is much broader than just hedge funds. Hedge fund compensation and CEO compensation are surprisingly similar, though the terminology is different. And the findings in Shue and Townsend are consistent with bad bets in “main street” firms too. Moreover, these compensation schemes can make a bad situation worse. Even transient

Figure 2
Actual Risk-taking and Performance in Response to Incentives



¹ Though redesigning contracts to fix this behavior without undoing good parts of the contracts is not obvious. But, since the financial crisis, it appears the industry is experimenting.

continued on flap...

bad luck can lead a manager to be below his threshold and thus take more risk and do worse, magnifying and perpetuating the harm—to the firm, the industry, or the economy. Beyond executives, millions of employees have stock options, sales people routinely have convex commission schemes, and others have performance pay schemes that add some convexity to their pay.

Risk taking is also a fundamental policy question. The beginnings of a policy response to risk taking, in the Dodd-Frank Act, focused on the link between compensation and incentives for risk taking. Yet, as the struggle around that legislation displayed, regulating incentives for risk taking is not simple, for a very fundamental reason. The same types of incentives that encourage valuable risk taking and effort can also induce inefficient and excessive risk taking.

Policy Implications

Since bad bets are bad both for society and for firms themselves,

public policymakers and firm decision makers are aligned and can respond by changing compensation structures and the jobs people do. From an incentive scheme design, two features can help—without important repercussions elsewhere. First, stop rewarding great outcomes that are unlikely to arise under sensible risk taking. These rewards provide little incentive for effort or sensible risks, but do encourage excessive risk taking. For example, the move toward non-stock bonus contracts can do this, if they are well designed. Second, find ways to expose managers to more of the downside of failure. The move to replace options with restricted stock does some of this. Finally, policymakers can encourage and firm decision makers can implement institutional risk controls, including empowering someone who is compensated to manage risk. The rise of the “chief risk officer” may serve this role, but only with the right authority and incentives. More generally, even in settings where the firm is risk tolerant, it is important to

design jobs and incentives with consideration for risk-taking incentives and controls.

References

de Figueriedo, Rui, Evan Rawley, and Orié Shelef. “Bad Bets: Excessive Risk Taking, Convex Incentives, and Performance.” SIEPR Discussion Paper 13-002, September 2013.

Huang, Jennifer, Clemens Sialm, and Hanjiang Zhang. “Risk Shifting and Mutual Fund Performance.” *Review of Financial Studies*, 24 (8), 2011, 2575-2616.

Palomino, Frédéric, and Andrea Prat. “Risk Taking and Optimal Contracts for Money Managers.” *The Rand Journal of Economics*, 34(1), 2003. 113-137.

Shue, Kelly, and Richard Townsend. “Swinging for the Fences: Executive Reactions to Quasi-Random Option Grants.” Chicago Booth Research Paper No. 13-03, September 2013.

SIEPR *policy brief*

Stanford University
366 Galvez Street
Stanford, CA 94305
MC 6015

A publication of the
Stanford Institute for
Economic Policy Research

Non-Profit Org.
U.S. Postage
PAID
Palo Alto, CA
Permit No. 28

SIEPR

About SIEPR

The Stanford Institute for Economic Policy Research (SIEPR) conducts research on important economic policy issues facing the United States and other countries. SIEPR's goal is to inform policymakers and to influence their decisions with long-term policy solutions.

Policy Briefs

SIEPR policy briefs are meant to inform and summarize important research by SIEPR faculty. Selecting a different economic topic each month, SIEPR will bring you up-to-date information and analysis on the issues involved.

SIEPR policy briefs reflect the views of the author. SIEPR is a non-partisan institute and does not take a stand on any issue.

For Additional Copies

Please see SIEPR website at
<http://SIEPR.stanford.edu>.