Creating Incentives for Innovation

Gustavo Manso

In an era of fast-paced technological change, innovation has become a business imperative. But it is not easy to make experimentation and risk-taking an integral part of an organization's business practice. This article describes a model for motivating innovation based on probability theory and experimental evidence. This research indicates that corporate leaders should create a culture that tolerates early failure and rewards long-term performance, and employees from the CEO down must be given incentives to innovate. These lessons are broadly applicable in fields beyond business such as scientific research and the academic world.

Businesses large and small find themselves at a crossroads as the pace of change accelerates. The rapid rise of technologies such as artificial intelligence, mobile networks, embedded sensors, and Big Data make it impossible to continue business as usual. In virtually every sector, technologically savvy new entrants are disrupting longstanding business models, forcing incumbents to adapt or die. In this environment, innovation is no longer simply desirable—it is an imperative that determines whether an organization will thrive or fall by the wayside.

The unforgiving demands of technological change are not lost on corporate leaders. Surveys repeatedly show they feel immense pressure to adopt new ways of doing business. For example, a 2016 KPMG survey of 400 U.S. chief executives found that keeping up with new technologies and responding as new competitors overturn business models ranked as the two

top concerns, cited by 81% and 76% of respondents respectively.[1] And in a 2017 PwC poll of 1379 CEOs in 79 countries, respondents said innovation was the aspect of their business they most wanted to strengthen.[2]

While it is easy to appreciate the importance of innovation, it is not simple to make it an integral part of an organization's business practice. Motivating managers to innovate remains a challenge for most businesses. Innovation requires the exploration of new, untested approaches that may very well fail, which flies in the face of the incentives that typically guide business behavior. Compensation arrangements are a key area. Standard pay-for-performance schemes that reward short-term financial results and punish failure with low compensation and possible termination may deter managers from trying new things in favor of tried-and-true methods that offer predictable outcomes. An organization that wants to encourage innovation must design compensation plans that free managers to take risks, experiment, and discover what practices and technologies are most effective. The essential elements of such plans are that early failure is tolerated, performance is measured over long time horizons, jobs are secure, and feedback is regular and timely.[3]

In the context of executive compensation, economic analysis indicates that the optimal innovation-motivating incentive scheme uses a combination of stock options with long vesting periods, option repricing, golden parachutes, and managerial entrenchment. With these provisions, compensation depends not only on total performance, but also on the path of performance. A manager who oversees a successful program of innovation may perform poorly at first but do well later, earning more than a conservative manager who does well initially but poorly later. Moreover, boards may be tempted to dismiss the innovating manager because of

poor early results, but entrenchment provides the time needed to allow new methods to bear fruit.

Provisions such as golden parachutes and option repricing are often frowned upon in the governance literature, and it is true that they are sometimes abused, rewarding managers despite poor business results. Nonetheless, they may be necessary parts of incentive schemes that motivate innovation. Restricting their use may discourage innovation at a time when it is vital for corporate survival.

More broadly, toleration of failure must be woven into organizational culture. Some of the business world's most successful and innovative companies, such as 3M and IBM, are known for giving their research departments great latitude, encouraging experimentation and risk-taking. Similarly, academic institutions grant faculty tenure, giving them license to pursue a research agenda without worrying about job security. And generous bankruptcy laws allow entrepreneurs to start over again when a business goes under. The larger point is that, if innovation is an organizational goal, then innovators must not be punished when their efforts fall short.

Incentive Schemes and Productivity

Since the emergence centuries ago of wage labor as way of organizing work, business owners have confronted the problem of how to motivate employees to be productive. Hourly wages or a fixed salary require workers only to make the minimum effort needed to keep their jobs. Academic literature typically analyses this problem using a principal-agent model. In these models, principals can create incentive plans that prompt agents to act in the principals' best

interest. These plans commonly take the form of pay-for-performance arrangements such as bonuses and piece work that offer monetary rewards for increased output.

Extensive research supports the idea that pay-for-performance boosts productivity. A variety of studies, including laboratory and field experiments, offer evidence that workers performing simple, routine tasks respond to financial incentives by exerting themselves more and performing better. In a study of agricultural worker in the Philippines, Foster and Rosenzweig (1994) show that, conditional on calorie intake, workers doing piece work lose more weight than workers on fixed wages, suggesting they work harder. Lazear (2000) shows that the productivity of Safelite Glass Corp. windshield installers increased when management changed their compensation from fixed wages to piece-rate pay. Shearer (2004) finds similar evidence in a randomized field experiment with Canadian tree planters. Ehrenberg and Bognanno (1990) document increased performance for professional golf players when monetary incentives are higher. Dickinson (1999) shows that subjects in a controlled laboratory experiment type more letters when their compensation is more sensitive to performance.

These studies are based on the standard principal-agent model, which focuses primarily on settings where the critical tension concerns whether the agent will work or shirk, and the intensity of the agent's effort is the sole determinant of output and productivity. This assumption may be appropriate for routine, repetitive tasks, but it may not apply when workers are confronted with open-ended, creative problems that call for innovative solutions. Incentives that encourage workers to install windshields at a faster pace may not be effective with employees who write software code, design marketing plans, teach a class, or manage a workforce.

Importantly, optimal performance may also require creativity and originality—in other words, innovation. Thus, compensation for open-ended, creative tasks calls for different kinds of incentive schemes. This is a pervasive problem. As routine, repetitive tasks are increasingly automated, a larger share of work is of the open-ended, creative type that requires new approaches and adaptability. It is important to consider what kinds of incentive schemes best fit the changing profile of work.

Modeling Innovation

This article considers a specific problem in crafting incentives plans for managers and employees whose jobs involve open-ended, creative tasks: how to encourage originality and innovation. To spur employees to try new things may require a different approach than what is needed to get them to do their jobs using familiar techniques.

To examine this question, I use a decision modeling method taken from probability theory called the bandit problem. This analytical approach was originally defined in terms of a gambler who must choose which slot machine (one-armed bandit) to play at a casino: Does she play the machine with the highest known reward or does she try several machines with the goal of finding those with even higher payoffs?[4] Innovation in this scenario is the discovery through experimentation and learning of courses of action that are superior to methods with well-known track records. In bandit problems, the agent is uncertain about the payoffs that different actions will generate. However, by trying new things, the agent can discover techniques that are more effective than known procedures. Thus, in bandit problems, there is a central tension between *exploration* and *exploitation*.[5] Exploration uncovers information about potentially

superior methods, but frequently leads down blind alleys, wasting time and effort. Exploitation relies on proven techniques that have high probabilities of reasonable payoff, but inhibit discovery of superior methods.

To study how organizations can create incentives for innovation, I build a model based on a bandit problem that has two participants: a principal and an agent. The principal wants to motivate the agent to engage in exploration. Many such situations arise in the business world: A corporate board may want to design a pro-innovation compensation package for a chief executive; similarly, a manager may want to encourage workers to experiment; a corporation may want to create a unit charged with developing new products; and a venture capital firm may want its portfolio companies to adopt potentially revolutionary, but unproven technologies. The question is, what incentives work best to get the agent to explore new methods?

My model has two periods and two possible outcomes in each period: success or failure. To focus on the tension between exploration and exploitation, I assume that in each period the agent may choose a conventional work method that has a known probability of success or a new work method with an unknown probability of success. Of course, the agent may also choose to shirk, that is, do nothing. But that doesn't raise the question of exploration versus exploitation. Instead, it involves the standard incentive problem of motivating the agent to be productive.

When the agent experiments with a new method in the first period, he is less likely to succeed than if he follows a known method. If the new method fails, he can go back to the conventional method in the second period. However, if the new method succeeds, then the

agent gains valuable information and may find that the innovation is better than the conventional method. Thus, if the agent discovers a successful new method, the total payoff from exploration may be higher than the payoff from exploitation. For that reason, the agent may be willing to try a new method even though the initial probability of success is lower.

In this model, the time horizon is critical. The discoveries arising from exploration become increasingly valuable if periods were added to the model. The more periods an agent is active, the higher the benefits of trying new techniques and the more willing the agent is to sacrifice payoffs early on by taking chances. As the time horizon expands, the agent has more opportunities to use the information he learns from experimentation.

Optimal Incentives for Innovation

The bandit problem model establishes two important points about innovation: the best results may involve initial failure and a substantial amount of time may be needed before the benefits of new techniques become apparent. These represent points of departure from standard incentive problems in the business world, supporting the hypothesis that incentive plans motivating exploration are fundamentally different from plans that motivate exploitation. Here we look at the key differences between the two types of reward systems.

Exploitation involves the repetition of well-known techniques. The optimal plan that motivates exploitation is based on the principle that the results of an agent's actions should be successful and predictable. In that regard, incentives for exploitation are similar to standard pay-for-performance arrangements used to motivate productivity in repeated tasks. The goal of such incentive plans is to encourage agents to perform well, which is measured in terms of

productivity or quality of work. From the point of view of principals, the main danger is that agents will shirk, that is, fail to carry out the tasks assigned to them.

On the other hand, since with exploration the agent is likely to waste time with unsuccessful actions, the optimal contract that motivates exploration has substantial tolerance—and may even reward—early failures. Moreover, since exploration uncovers information that is useful for future decisions, the optimal contract that motivates exploration rewards long-term success.

Termination

The threat of termination raises a special problem. If we believe managers should be accountable for the results of their work, then getting fired represents the ultimate sanction when performance is unacceptable. But what is acceptable and what is unacceptable may be very different when the agent is expected to engage in exploration as opposed to exploitation. The threat of termination discourages agents from shirking or exploring new actions, and thus motivates exploitation. In fact, an aggressive approach to termination may be optimal if principals want to get the best possible results from known business practices.

By contrast, the effects of termination on incentives for exploration are ambiguous. To be sure, fear of getting fired prevents an agent from shirking. At the same time though, the danger of job loss encourages the agent to take a safe course by acting conventionally and avoiding the unknown. If principals are primarily worried that the agent won't work hard, they may want to have a credible threat of termination. But if their main concern is motivating the agent to innovate, they need to provide assurances that the agent's jobs is secure. In summary,

termination is useful to prevent shirking, but may be a disincentive to innovation if it encourages the agent to fall back on familiar ways of doing business.

Commitment to Long-Term Contracts

Similarly, the ideal contract that motivates innovation focuses on results over an extended period, rewarding long-term, but not short-term, success. It may even reward short-term failure. The key intuition is that if an agent is not protected against early failure, she may prefer to exploit conventional work methods. Offering the agent a safety net allows her to venture into the unknown knowing she will have sufficient time to develop her idea. This even holds if an innovation shows early success. Additional time allows the agent to adjust and improve a new work method, generating increasingly effective results as experience is gained. For this reason, delaying compensation while innovations are perfected is optimal for principals. The willingness of principals to commit to a long-term contract is critical if they want agents to explore rather than exploit.

Explicit contracts and corporate culture

What does motivating innovation mean in practice? Two key areas to look at are employment agreements with executives and senior managers, and the corporate cultural environment. In the case of executive compensation, companies can encourage innovation by offering explicit contracts that include stock options with long vesting periods, option repricing, golden parachutes, and entrenchment. These provisions reward long-term results and provide job security.

However, such employment agreements may not be practical for lower-level employees. Large corporations often claim that it is hard to get rank-and-file staff members to

be more creative. One reason may be employee cynicism. Employees may doubt that new thinking is genuinely appreciated, their innovations will be rewarded, and early failures tolerated, no matter how ringing the pro-innovation rhetoric may be from on high. This is where corporate culture can make a crucial difference. A culture that fosters experimentation, rewards innovation, and gives employees freedom and time to develop new ideas is essential to overcome the innate caution and skepticism of the workforce.[6] Unlike explicit contracts, corporate culture is intangible. A corporation gains credibility and effectively motivates employees to take risks only to the extent that it builds a reputation for rewarding innovation. That reputation must be earned—the company's deeds must match its words. It takes time and follow-through to win employee trust.

Although culture plays the key role, there are nevertheless explicit tools for motivating innovation below the executive level. Corporate and academic research departments often grant tenure to researchers, which frees them to explore ideas that will probably fail but could lead to breakthroughs. In addition, researchers can be rewarded in a variety of ways for long-term success. Research shows that when managers in charge of research and development get more long-term incentives, their patents are more heavily cited.[7]

Feedback

A large academic literature focuses on the importance of feedback for business performance. For example, focus groups and other consumer tests generate information on market attitudes about a company's products and services. And suppliers may learn about new technologies by communicating with their customers.[8] In the principal-agent relationship, principals frequently have knowledge that agents lack and may be better positioned to evaluate

performance. A venture capitalist may know more about a market niche and an enterprise's commercial value than the entrepreneur being funded. And company executives may have superior information about consumer attitudes and the market potential of new products than the company's research and development personnel. This raises the question of whether principles should communicate their insights to agents to help them improve performance. In my model, the answer depends on whether the principals want to motivate exploitation or exploration. If exploitation is desired, then avoiding feedback on performance is optimal. But feedback is essential if the goal is to get agents to explore. In short, principals should never provide feedback on performance to the agent when implementing exploitation, but should always provide feedback when implementing exploration.

The key point here concerns adjustments to performance. By definition, exploitation involves repetition of known actions. In other words, principals seeking exploitation do not want an agent to make adjustments—they simply want her to do what she has done successfully in the past. Therefore, the principals do not communicate their observations on performance to the agent. This result is similar to previous research, which has found that when principals want to induce an agent to carry out a repetitive effort, they should not communicate information about performance to the agent.[9]

The situation is completely different if the principals want the agent to explore new methods. Exploration requires interim adjustments to action based on performance. The principals may have insights that will help make the agent's experimentation more efficient. Regular feedback is needed if the goal is to motivate exploration and principals are better equipped to evaluate performance than the agent. The purpose of feedback is to provide

information that guides adjustments and improves the agent's performance. This is not the same as a performance review used to make compensation and promotion decisions. No rewards and punishments are linked to the feedback discussed here. It is purely informational in that it communicates observations that may help an agent achieve long-term success in an innovation agenda.

Bankruptcy Laws

The idea that toleration of failure fosters innovation applies to entrepreneurs as well as employees. If someone starting a business borrows funds to support an unproven product or strategy and the innovation fails to make money, then the business owner may not be able to pay her debts. She will be insolvent and her prime recourse may be bankruptcy. Innovation-friendly bankruptcy laws—those that encourage exploration—will discharge the entrepreneur's debt and give her a fresh start, allowing her to undertake new projects. The failure of her original business does not have dire economic consequences and she is free to initiate another risk-taking project. Some research suggests that such debtor friendly bankruptcy laws lead to more innovation.[10]

The European Council explicitly recognized this principle in its "European Charter for Small Enterprises" in 2000. The charter states that "failure is concomitant with responsible initiative and risk-taking and must be mainly envisaged as a learning opportunity." It called for bankruptcy law reforms that would permit failed entrepreneurs to start over again. The United

States has gone in the other direction, adopting a creditor-friendly bankruptcy law in 2005 that makes it harder for entrepreneurs to discharge their debts after failure.

How Independent Boards and Shareholder Litigation Influence Innovation

Independent Boards

An independent board of directors is another area where widely accepted corporate governance principles and optimal incentives for innovation may be at odds, although data on this question are ambiguous. An extensive literature argues that a majority of directors should have no ties to a company other than their board seats. Independence allows them to ride herd on managers, ensuring that executives act in the interests of shareholders and sanctioning them if performance is unacceptable. But how an independent board influences innovation is less clear. If independent directors put pressure on executives to improve short-term financial results and avoid taking risk, they may be creating disincentives for innovation.

Benjamin Balsmeier of the Swiss Economic Institute, Lee Fleming of UC Berkeley, and I investigated these effects by reviewing companies that were forced by legislative and regulatory changes to make their boards more independent.[11] Starting in 2002, stock exchanges and the Sarbanes-Oxley Act required that a majority of public company directors be independent. By studying patent and citation patterns, we found that the impact of shifting to a more independent board was large. Innovation increased substantially, but we identified no significant effect on more speculative and exploratory forms of innovation. Instead, companies focused on familiar areas of technology in which their competitors were also active.

Overall, patents increased between 20% and 30%, but those were generally for incremental technological advances, not breakthroughs. Patent citations rose 40% to 60%, a significant increase. However, when we looked at citation distribution, we found no change in the numbers of uncited and highly cited patents, which may be associated with more radical innovation strategies. The story appears to be that strengthened board oversight makes executives more productive as measured by number of patents, but these patents are concentrated in existing technologies, not in more experimental, cutting-edge areas of research. Heightened board pressure may be encouraging managers to exploit already successful areas of expertise at the expense of riskier avenues of exploration.

Shareholder Litigation

An important hazard faced by managers of public companies is the threat of lawsuits accusing them of breaching their fiduciary responsibilities to shareholders. A large legal industry is devoted to seeking out instances in which it can be argued that managers acted in their personal best interests at the expense of shareholders. Undoubtedly, this legal liability disciplines managers by raising the costs of negligence and self-dealing. At the same time though, theory suggests that fear of shareholder litigation could stifle innovation. Managers may have strong motives to reduce their legal exposure, which could feed short-term thinking and risk aversion. Thus, risk of lawsuits may impair both the quantity and quality of innovation. Innovation involves risk. But, fearing liability, companies may play it safe by cutting back on exploration and concentrating on what is legally safe, but not groundbreaking.

Chen Lin and Sibo Liu, both of the University of Hong Kong, and I looked for empirical evidence on this hypothesis by studying adoption of universal demand (UD) laws in 23 U.S.

states from 1989 to 2005.[12] UD laws regulate derivative lawsuits in which a shareholder sues a third party—typically a company insider—on behalf of a corporation. These laws create an obstacle to litigation by requiring shareholders to demand that a company's board of directors take remedial action before they file suit. In states that have passed UD laws, the incidence of derivative lawsuits has fallen.

The staggered adoption of UD laws allows us to study the relationship between shareholder litigation and corporate innovation. Using a sample consisting of 4,506 U.S. public companies firms between 1976 and 2006, we find evidence that UD laws encouraged innovation by reducing the threat of shareholder litigation. Following passage of these laws, companies boosted research and development spending by about 11% of the sample mean for R&D relative to companies in states without such laws. Average citations per patent rose and a greater number of patents generated large numbers of citations, indications of improved patent quality. Moreover, more patents were in experimental and unfamiliar technological areas, a sign that companies were doing more exploration. Taken together, these results appear to confirm the hypothesis that shareholder litigation impedes corporate innovation

Experimental Evidence on Motivating Innovation

The argument that optimal incentives for innovation are fundamentally different from those for exploitation is based on a mathematical model that uses the bandit problem to examine the business relationship between principals and agents. To test this idea, Florian Ederer, now of the Yale School of Management, and I conducted a laboratory experiment to investigate the tension between use of known methods and exploration.[13] Specifically, we

looked for evidence that tolerance for early failure and reward for long-term success motivates innovation. We recruited 379 subjects to operate a computerized lemonade stand and gave them a choice between making minor adjustments to the business decisions made by a previous manager or carrying out major changes to discover a better strategy, such as moving to a different location or radically altering the product mix. The experiment was designed in such a way that a certain set of product and location choices represented the optimal business strategy.

The subjects were divided into three groups with different compensation schemes. The first group received a fixed wage in each period of the experiment; the second got a standard pay-for-performance contract allotting them a fixed percentage of profits produced during the experiment; the third received contracts designed to motivate exploration, basing their compensation on a fixed percentage of profits generated in the second half of the experiment. The third contract included two provisions designed to encourage participants to explore. First, subjects could fail at no cost in the first half of the experiment while they experimented with new business strategies. Second, they were rewarded for their performance at the end of the experiment, which gave them time to learn better ways of operating the lemonade stand. Our goal was to see which incentive regime produced superior business strategies. We hypothesized that subjects under exploration contracts were more likely to experiment and find improved strategies than those working under fixed-wage or pay-for-performance contracts.

Our results strongly supported the idea that toleration of early failure combined with long-term incentives encourage innovation. By the end of the experiment, subjects under the exploration contract found the best location for the lemonade stand 80% of the time compared

with 60% and 40% respectively for those under the fixed-wage and the pay-for-performance contracts. Detailed analysis of the results provided important insights. Subjects under the fixed-wage contract did a significant amount of exploration, but they did not go about it as systematically as their counterparts under an exploration contract. When we examined the notes subjects took in a table provided to them at the beginning of the experiment, we found that only 55% of those under the fixed-wage contract regularly used it to track their business decisions and profits. By contrast, 82% of those under an exploration contract used the table to monitor their operations.

We also found that subjects under a pay-for-performance contract tended to put most of their energy into fine-tuning the previous manager's product mix instead of searching for better locations. During the experiment's first 10 periods, 80% of those under the exploration contract left the original location for another place to do business; subjects under the pay-for-performance contract relocated only 50% of the time. Risk aversion played an important role in explaining differences in behavior and performance of subjects under exploration and pay-for-performance contracts. Under pay-for-performance contracts, risk-averse subjects were less likely to discover the optimal strategy and generally posted lower profits than their less-risk-averse counterparts.

To study the effects of termination, we introduced two new groups into the experiment: a regular termination group and a golden parachute termination group. Subjects in both groups got exploration contracts and were advised the experiment would end early if their profits in the first 10 periods fell below a certain threshold. However, those in the golden parachute termination group were also told that they would get a special payment if the experiment

ended early. We found that 65% of subjects in the golden parachute termination group discovered the best business location compared with only 45% of those in the termination group without golden parachutes, suggesting the promise of payment in the event of failure encouraged subjects to take chances.

In summary, the experiment showed that subjects under exploration incentive schemes that tolerate early failure and reward long-term performance experiment more and are more likely to discover novel business strategies than those under fixed-wage and standard pay-for-performance incentive schemes. In addition, the threat of termination can undermine incentives for innovation, but golden parachutes mitigate those innovation-reducing effects.

Incentives for Exploration in Scientific Research

So far, this article has examined incentives for innovation in business settings. Does the fundamental insight that patience, toleration of early failure, and focus on long-term results encourage innovation apply in other areas as well? One field to look at is scientific research, which takes place under a variety of regimes. Some are forgiving, offering researchers ample time to explore different avenues and develop their ideas; others subject researchers to intense pressure to produce results quickly. How do these varying approaches affect scientific innovation?

Pierre Azoulay of MIT, Joshua Graff Zivin of UC San Diego, and I examined this question by comparing researchers at the Howard Hughes Medical Institute (HHMI) with National Institutes of Health (NIH) grantees.[14] These two institutions offer a superb test of the effects of a long time horizon and toleration of early failure on innovation. HHMI is known for rewarding

long-term success and giving investigators great freedom to experiment. The value it places on creativity could hardly be more explicit. Its website urges investigators "to take risks, to explore unproven avenues, to embrace the unknown—even if it means uncertainty or the chance of failure."[15] HHMI award cycles are five years. Initial reviews are forgiving, while later reviews provide researchers detailed feedback from distinguished panels of scientists. In addition, HHMI has a policy that is notably effective in fostering exploration—it selects "people, not projects." Under this standard, early failures don't trigger sanctions. That gives researchers more latitude, enabling them to easily shift resources to new lines of investigation if the initial ones don't pan out.

By contrast, life scientists funded by the NIH typically receive three-year grants. Renewal is jeopardized by failure and feedback on performance is limited. Importantly, grants are earmarked for specific projects, not individual scientists, which discourages recipients from pursuing lines of research other than those explicitly funded by the NIH. Reviews are more cumbersome than those of HHMI and grantees must demonstrate in detail that they have accomplished what they promised in their grant applications.

TABLE 1 Comparison between the Two Sources of Funding

NIH R01 Grants	HHMI Investigator Program
Three- to five-year funding	Five-year funding
First review is similar to any other review	First review is rather lax
Funds dry up upon nonrenewal	Two-year phase-down upon nonrenewal
Some feedback in the renewal process	Feedback from renowned scientists
Funding is for a particular project	"People, not projects"

The concepts of exploitation and exploration have different connotations in science than in business. In our examination of research in the life sciences, we defined incremental

work that fine-tunes known approaches as a proxy for exploitation and more radical avenues of investigation that break new ground as a proxy for exploration. Our results show that incentives that encourage innovation in business have markedly similar effects in science. To measure exploration, we used citations in the scientific literature, reasoning that the more pathbreaking the research, the more likely it was to be cited. We found that HHMI investigators produced high-impact articles at a much higher rate than a control group of similarly accomplished NIH-funded scientists. Participation in the HHMI program increased overall publication output by 39%, but that figure jumped to 96% when we looked at the number of publications in the top percentile of citations. At the same time, we found that HHMI-supported scientists "flopped" more often than the NIH researchers: compared with the control group, 35% more of their articles were cited less frequently than their own least-well-cited pre-appointment publication. This is evidence that HHMI investigators were taking more chances, that is, they were engaging in more exploration, as our theory suggests they would.

Overall, we found that the HHMI investigator program, which rewards long-term success, encourages experimentation, and offers extensive feedback to grantees, leads to more breakthroughs than NIH funding, which has shorter grant cycles, predefined research agendas, and stricter review policies. These results provide empirical support from outside the business world to our findings on incentives for innovation.

Conclusion

Over the past several decades, a series of high-profile corporate scandals at companies such as Enron and WorldCom Communications have cast a spotlight on governance issues and

fueled a reform movement aimed at ensuring the interests of managers and shareholders are better aligned. The central principles of this movement are that managers must be accountable to shareholders and that compensation be closely linked to performance. Furthermore, good governance implies that the use of incentives such as stock options, option repricing, and golden parachutes not be abused. Yet, even as the pay-for-performance movement has gathered steam, other voices have been heard decrying the prevalence of short-term thinking in the corporate world. Scholars, public officials, and business leaders like Warren Buffett have warned that long-term value creation is getting short shrift as managers concentrate on pumping up quarterly results.

It may be that these two arguments for reform are at least partially contradictory. Demands that managers avoid self-dealing and get compensated based on performance are laudable. They arise from widespread misconduct that has weakened corporate performance and victimized shareholders. Nonetheless, if short-term results serve as measurements of performance, then these reforms carry a cost. To generate good quarterly earnings, managers may avoid risk-taking and exploration—precisely the steps that can multiply long-term value. If innovation is a goal, then managers must be allowed to fail and must be given time for new methods to prove their worth. Thus, incentive plans should include options with long vesting periods and allow them to reprice if share prices fall while an innovation agenda is pursued. Managers should not lose their jobs if short-term results suffer. And, if they do leave, golden parachutes should pay them for the exploration they supervised.

It may not be easy to balance the competing demands of good governance and innovation. It increases the burden on corporate directors, who are now called on to carry out a

two-pronged oversight agenda: they must make sure managers are acting in shareholders' interests and, at the same time, create incentives for innovation. That requires them to be vigilant and patient at the same time. Judging managerial performance becomes much harder if short-term results are not used as the yardstick. In addition, demands on board members to communicate with managers intensify because feedback on performance is crucial for supporting innovation. Nonetheless, if innovation is now the greatest challenge in the business world, then it only makes sense to give corporate managers the right incentives to break new ground.

Innovation involves more than the terms of compensation agreements and employment contracts. To encourage experimentation and risk-taking, managers must consider the culture of their organizations. Does that culture give people the time and space to carry out projects that may not produce immediate rewards? Does it stigmatize those who fail or does it reward those who take chances? A model for encouraging innovation may be found at some research-oriented universities. At the most innovative institutions, publish-or-perish rules may be relaxed and tenure granted to scholars whose work holds the promise of breaking new ground.

This article has focused on management questions, but the point should also be stressed that public policy also has an important role to play in providing incentives for innovation. A voluminous literature examines this subject. One example considered here is bankruptcy laws: Such laws can discourage risk-taking by punishing businesses that fail or they can support innovation by giving those whose projects don't pan out the opportunity to try again. A society that wants to bolster its position in an increasingly competitive world must make promotion of innovation a central goal of policy.

[1] KPMG, "Now or Never: CEOs Mobilize for the Fourth Industrial Revolution," U.S. CEO Outlook 2016, https://assets.kpmg.com/content/dam/kpmg/pdf/2016/07/2016-ceo-survey.pdf.
[2] PWC, "20 Years Inside the Mind of the CEO… What's Next?" 20th CEO Survey, 2017, https://www.pwc.com/gx/en/ceo-survey/2017/pwc-ceo-20th-survey-report-2017.pdf.
[3] G. Manso, "Motivating Innovation," *The Journal of Finance*, 56/5 (October 2011): 1823-1860.
[4] D. Berry and B. Fristedt, *Bandit Problems: Sequential Allocation of Experiments, Monographs on Statistics and Applied Probability* (London: Chapman and Hall, 1985).
[5] J.G. March, "Exploration and exploitation in organizational learning," *Organization Science,* 2/1 (February 1991): 71–87.
[6] R. Farson and R. Keyes, *Whoever Makes the Most Mistakes Wins: The Paradox of Innovation* (New York, NY: Free Press, 2002) and R. Sutton, *Weird Ideas that Work: 11 1/2 Practices for Promoting, Managing, and Sustaining Innovation* (New York, NY: Free Press, 2002) provide examples of innovative corporations.
[7] J. Lerner and J. Wulf, "Innovation and Incentives: Evidence from Corporate R&D," *The Review of Economics and Statistics,* 89/4 (November 2007): 634–644.
[8] Y. Chu, X. Tian, and W. Wang, "Learning from Customers: Corporate Innovation along the Supply Chain," AFA 2015 Boston Meetings (Aug. 8, 2015), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2505029.
[9] A. Lizzeri, M. Meyer, and N. Persico, "The Incentive Effects of Interim Performance Evaluations," Working paper, CARESS, 2002. W. Fuchs, "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review*, 97/4 (September 2007): 1432–1448.
[10] V. Acharya and K. Subramanian, "Bankruptcy Codes and Innovation," *The Review of Financial Studies, 22/12* (December 2009): 4949-4988.
[11] B. Balsmeier,, L. Fleming, and G. Manso, "Independent Boards and Innovation," *Journal of Financial Economics*, 123/3 (March 2017): 536–57.
[12] C. Lin, S. Liu, and G. Manso, "Shareholder Litigation and Corporate Innovation," unpublished manuscript.
[13] F. Ederer and G. Manso, "Is Pay for Performance Detrimental to Innovation?" *Management Science*, 59/7 (February 2013): 1496-1513.
[14] P. Azoulay, J. Graff Zivin, and G. Manso, "Incentives and Creativity: Evidence from the Academic Life Sciences," *RAND Journal of Economics*, 42/3 (Fall 2011): 527–554.
[15] See http://www.hhmi.org/programs/biomedical-research/investigator-program.