



# Understanding Corporate Governance Through Learning Models of Managerial Competence\*

Benjamin E. Hermalin\*\*

University of California, Berkeley, United States and ECGI, Brussels, Belgium

Michael S. Weisbach

Ohio State University and NBER, United States and ECGI, Brussels, Belgium

Received 29 November 2018; Accepted 28 December 2018

## Abstract

Shareholders, the board of directors, and potential future employers are continually assessing any CEO's quality. As documented by an extensive literature, such assessment plays a critical role in corporate governance because it generates incentives (good and bad), introduces assorted risks, and affects the various battles that rage among the relevant actors for corporate control. Consequently, assessment (or learning) is a key perspective from which to study, understand, and possibly even regulate corporate governance. Moreover, because learning is a behavior notoriously subject to systematic biases, assessment is a natural avenue through which to introduce behavioral and psychological insights into the study of corporate governance.

**Keywords** Corporate governance; Career concerns; Learning and assessment; Cognitive biases

*JEL Classification:* G34, M12, D83, D81

## 1. Introduction

Within economics, corporate governance has long been viewed through the lens of agency: the fear that a firm's managers are inclined to behave at odds with the

---

\*This paper is a shorter and less technical discussion of the issues addressed in Chapter 3 of *The Handbook of the Economics of Corporate Governance* (Hermalin and Weisbach, 2017a) and draws heavily from it. The authors thank Jongha Lim, Yihui Pan, Miriam Schwartz-Ziv, Berk Sensoy, Luke Taylor, Ralph Walkling, and Tracy Wang for helpful comments on an earlier draft, and Shan Ge for excellent research assistance. Hermalin gratefully acknowledges the financial support of the Thomas & Alison Schneider Distinguished Professorship in Finance.

\*\*Corresponding author: University of California, Berkeley, CA, United States. Tel: +1-510-642-6474, email: hermalin@berkeley.edu.

desires of its principals (typically, the firm's shareholders).<sup>1</sup> Yet, without minimizing the importance of agency, there is another key aspect to governance that has received much less attention, namely the way in which managers are assessed with respect to their intrinsic ability, competency, and match with their employer, and the effects such assessments have on the behavior of the relevant actors.

A long recognized scientific principle is that observation or measurement can directly affect the phenomenon being observed or measured. In a governance context, this principle manifests in the following way: because they know that others seek to make inferences about them, managers know there are personal consequences of being observed, an insight attributable to Fama (1980) and Holmstrom (1982).<sup>2</sup> Furthermore, this insight has important implications for incentive provision and corporate governance. In particular, because the outcome of an assessment, which can affect a manager's employment, pay, or both, is unknown in advance, the act of observation necessarily exposes the manager to risk. Consequently, in their contracts, managers will demand compensation for this risk, so management compensation will be a function of the way in which the assessment is carried out. Furthermore, since managers are strategic players, they can, therefore, be expected to influence how they will be observed and by whom. Assessment should, therefore, affect corporate governance by generating incentives (some good, some bad); by creating tradeoffs between how accurately managers are assessed and the risk thereby imposed on them; and by motivating managers to seek to influence by whom they are observed (i.e., governed).

The way in which the relevant parties learn about managerial ability can explain a number of factors related to firms' governance. First, the managerial labor market's assessment of a manager's ability can provide him incentives to perform well, because such assessment affects his future wages. However, such incentives are unlikely to be optimal, often being weaker than ideal, but, in other instances, being too strong. Worse, they can potentially have perverse effects, such as distorting managers' decisions about investment or project choice. Second, assessment is a critical part of the process by which managers are chosen and fired. Consequently, management turnover is naturally studied through a learning framework. Third, what *will be* learned is unknown *ex ante*; hence,

---

<sup>1</sup>Smith (1776, p. 700): "The directors of [joint stock] companies, however, being the managers rather of other people's money than of their own, it cannot well be expected, that they should watch over it with the same anxious vigilance [as owners] . . . Negligence and profusion, therefore, must always prevail, more or less, in the management of the affairs of such a company." In the modern era, there was a revived interest in agency problems and their relation to governance led by Berle and Means (1932), Williamson (1963), and Jensen and Meckling (1976), among others. (See Becht *et al.* (2003) and Hermalin (2013) for recent surveys of the relevant literature as well as Hermalin and Weisbach (2017b).)

<sup>2</sup>Holmstrom's paper was originally published in 1982, in a hard to find festschrift for Lars Wahlbeck. In the pre-PDF 1980s, a photocopy of the working-paper version of the paper was a treasured belonging of students and scholars interested in contract theory and governance. In 1999, the *Review of Economic Studies* wisely reprinted the paper (Holmstrom, 1999).

things that are functions of what is learned, such as future compensation and job retention, are necessarily uncertain *ex ante*, which means learning and assessment expose managers to risk. Exposure to such risk will affect their compensation demands, as well as their behavior more generally. Fourth, a manager judged to be better than average becomes a “rare commodity,” which gives him or her bargaining power. This too affects compensation; but it also puts the manager in position to bargain for more control over who governs him or her (e.g., who sits on the board of directors), leading to problems of entrenchment. Finally, learning about the ability of a firm’s top management will cause a rational stock market to update its assessment of its future profits; hence, systematic relationships should exist between learning and stock return volatility. In this article, we draw from the theoretical and empirical literatures to elaborate many of these points, and argue that they are of first-order importance for both normative and positive analyses of governance. A more complete, albeit also more technical, survey can be found in Hermalin and Weisbach (2017a).

It is worth noting that the effects enumerated above could be even more important if the assessors are “bad Bayesians” and fall foul to any of a number of well-documented cognitive biases (e.g., underweighting base rates, attribution bias, and so forth). Underweighting base rates, for instance, means that observers overweight a manager’s current performance relative to his or her past record, which makes observers’ responses to current performance more volatile than a rational-actor model might predict. Or, for instance, attribution bias means that observers tend to place too much importance on managers’ roles in affecting success or failure, which makes these observers overly sensitive to what they are learning. In short, as we discuss at the end of this paper, an important way in which behavioral insights could matter for corporate governance is through their implications for how managers’ abilities are assessed. Further, while behavioral factors can sometimes mute the insights from standard agency theory (e.g., intrinsic motivations lessening the importance of extrinsic incentives), behavioral biases, in particular those that make people bad Bayesians, will tend to amplify the consequences of learning (e.g., making actors “learn too much”).

## 2. The Good, the Bad, and the Ugly

Consider a competitive market for managerial talent. In such a market, the more able a manager appears, the greater will be the compensation she can command (the higher will be the bids of competing firms for her services). Ability is inferred from performance; hence, the more successful is the firm one currently manages, the greater will be one’s future compensation.<sup>3</sup> Performance is not, though, only a function of innate ability: it can also depend on the actions the manager takes. Consequently, as originally suggested by Fama (1980), because of their effect on the

---

<sup>3</sup>This is true even if the manager stays with her current employer: the current employer will need to increase the manager’s compensation in response to the outside offers her good performance elicits.

market's inference about her ability and, thus, on her compensation, the manager has incentives to take actions that will cause her firm to perform well.

It is worth making three points about these good incentives induced by assessment. First, as Holmstrom (1982) shows, such incentives are generically *sub-optimal*: they can be weaker than ideal (induce less than first-best actions) or stronger than ideal (induce greater than first-best actions), but generally will fail to be just right. Second, as Fama acknowledges and Holmstrom demonstrates formally, these incentives tend not be constant over the course of the manager's career: the more of a known entity she becomes, the less influence any one period's performance will have on estimates of her ability (at least ought to have if decision makers are Bayesian), and, thus, the less powerful are her incentives to take actions the shareholders desire. Third, there is an important game-theoretic subtlety to note: although the manager pursues actions to affect the market's assessment of her ability, the actors in that market anticipate her incentive to do so and, thus, "subtract" the actions they expect her to take from realized performance before forming their estimate of her ability.

Another way of stating this third point is that, in equilibrium, no one is fooled. One might, therefore, ask why the manager tries to influence the market's estimate through her choice of actions if the market will see through that. The answer is that the market subtracts the actions it *expects* her to take—it does not see what she actually chooses to do—so if she chose not to take those expected actions, then her performance would be worse and the market would subtract its expectation from that worse performance, so her inferred ability would be correspondingly low. In essence, the manager is like the Red Queen in Lewis Carroll's *Through the Looking Glass*: she must run as fast as possible just to stay still.

A bit of formalism will help flesh out these ideas. Suppose the value (gross return to her employer) that a manager generates in a given period is  $e + x$ , where  $e$  is her effort and  $x$  is a normally distributed random variable with a mean equal to the manager's ability,  $\alpha$ , and known variance. Likewise, the manager's ability is drawn from a normal distribution with known mean and variance.<sup>4</sup> Assume that the manager's per-period utility is  $w - c(e)$ , where  $w$  is her compensation and  $c : \mathcal{R}_+ \rightarrow \mathcal{R}_+$  is her cost-of-effort function.<sup>5</sup> Compensation,  $w$ , is set each period by a competitive market for the manager so that it equals the expected value she will produce in equilibrium.<sup>6</sup> Observe that these

---

<sup>4</sup>The literature has used the normal distribution because it has desirable properties for studying learning: updating rules are straightforward and all posterior distributions are also normal.

<sup>5</sup>In what follows, it is to be understood that  $c(\cdot)$  satisfies the "standard" assumptions that ensure solutions exist to the optimization programs to be considered, that these solutions are interior, and that they vary continuously with respect to the exogenous parameters.

<sup>6</sup>Note that there is a subtle assumption that compensation is based on the value the manager is expected to generate, not the value she does. Such an assumption can be justified in a number of ways (see, e.g., Holmstrom, 1982). The principal reason to make this assumption is in order to contrast reputational incentives with formal incentives.

assumptions imply that actual managerial effort in a given period does not affect the manager's compensation in that period. Therefore, a manager has no incentive to supply any effort in the last period of her career—if there is no future, there are no reputational or career concerns, a dramatic illustration of the general principal that reputational incentives necessarily “die out” toward the end of a career.

Consider a manager with a two-period career. As just noted, in the second (last) period, she will supply no effort. Hence, her second-period compensation will be the expected value of  $x$  (recall competition for her services drives  $w$  to equal the value of what she is expected to produce). The expected value of  $x$ , in turn, equals the posterior estimate of her ability,  $\hat{\alpha}$  (i.e., the estimate based on prior beliefs and first-period performance). It can be shown (see Holmstrom, 1982 for details) that:

$$\hat{\alpha} = \frac{\eta}{\eta + \rho}(x + e - e^E) + \frac{\rho}{\eta + \rho}\alpha_0, \quad (1)$$

where  $1/\eta$  is the variance of  $x$ ,<sup>7</sup>  $1/\rho$  is the variance of  $\alpha$ ,  $e^E$  is the effort that others *anticipate* the manager will expend (in equilibrium, they anticipate correctly; this is the Red Queen effect), and  $\alpha_0$  is the prior estimate of ability (in this case, the mean of ability). In the first period, the manager will choose  $e$  to maximize the difference between its expected benefit to her future compensation and her current cost of effort of  $\hat{\alpha} - c(e)$ .<sup>8</sup> Her marginal return to effort equals:

$$\frac{\eta}{\eta + \rho} < 1.$$

Given that the *social* marginal return is 1, it follows that her incentives are too weak: the manager supplies less than the socially optimal (first-best) level of effort. This conclusion, it needs be noted, does not necessarily carry over to a model that is more than two periods. While no effort is supplied in the last period (as noted above) and there is less than optimal effort in the penultimate period, there is the confounding effect, when there are more than two periods, that, because the earliest periods of her career influence compensation over *all* future periods, the manager's effort incentives are exceedingly strong; hence, early in her career, the manager's incentives could exceed the social optimal (see Holmstrom, 1982, or §3.1 of Hermlin and Weisbach, 2017a). What, however, is true, even in a more general multi-period model, is that the manager's private incentives (desire to maximize expected compensation less her cost of effort) will not generically correspond to what is socially optimal.

<sup>7</sup>The multiplicative inverse of the variance (e.g.,  $\eta$ ) is known in the literature as the *precision* of the distribution.

<sup>8</sup>For convenience, we ignore discounting (i.e., the fact that the cost is borne today, but the benefit realized in the future).

The discussion to this point shows that while reputational concerns are unlikely to create perfect incentives, they do motivate managers to some extent. However, one must be careful to consider what they are motivating managers to do. To this point, the assumption has been that managerial effort is socially beneficial. But reputational or career concerns can also incentivize behavior that is not socially desirable. Indeed, there are models in which career concerns create incentives that work against value maximization.

To illustrate one way in which reputational incentives can be counterproductive, suppose, now, that what the manager does is to decide whether or not to pursue a project (note, in this model variant there is no effort; i.e., no  $e$  nor  $e^E$ ). Whether she has a project to pursue is uncertain and her private information. If she undertakes a project, it pays off,  $x$ . Suppose the prior estimate of her ability,  $\alpha_0$ , which is also the expected value of  $x$ , is positive; consequently, if she has a project, it should be undertaken given it has a positive expected value. As in the previous model, her second-period compensation is the posterior estimate of her ability,  $\hat{\alpha}$ . The posterior estimate,  $\hat{\alpha}$ , is given by equation (1) if a project is undertaken. If no project is pursued the posterior estimate is necessarily the prior estimate,  $\alpha_0$  (given there would then be no new information on which to revise that estimate). Finally, suppose, now, the manager is risk averse in future income. It follows from (1) that undertaking the project exposes the manager to risk without affecting her expected compensation (the expectation of (1) with  $e \equiv e^E \equiv 0$  is, as noted,  $\alpha_0$ ). Given her risk aversion, it follows that she will never undertake a project should she have one. In other words, as Holmstrom (1982) originally showed, career concerns can create the wrong incentives when it comes to undertaking risky projects.<sup>9</sup>

In addition to distorting decisions about whether to pursue projects (more generally, which projects to take), reputational effects also distort intertemporal trade-offs. Consider a third model in which, now,  $e$  represents efforts that boost short-term (current-period) profits, *but at the expense of future profits*, and in such a way

---

<sup>9</sup>Hermalin (1993) shows that if the manager must pursue some project, but projects vary in terms of risk and the risk (i.e.,  $1/\eta$ ) is observable, then the manager can have incentives to pursue the riskiest projects, regardless of expected payoff: the greater is  $1/\eta$ , the smaller is  $\eta$ . From (1), this means more weight is being put on the non-stochastic element and less on the stochastic element; that is, the riskier the project chosen, the less the risk to the manager's compensation. Hence, career concerns can lead both to excess caution (eschewing projects), as in Holmstrom (1982) and Holmstrom and Ricart i Costa (1986), or to taking on excessive risk, as in Hermalin (1993). (See the last for more on this issue.) Other papers that explore the relation between project choice and career concerns include Scharfstein and Stein (1990), Hirshleifer and Thakor (1992), and Milbourn *et al.* (2001). Dominguez-Martinez *et al.* (2008) provide a variant on these models in which the CEO knows his ability (in contrast to the usual assumption of the literature). Now the choice to undertake a project becomes a signal of ability, which affects a number of the conclusions (in particular, high-ability CEOs may be "overactive" and boards "lenient" in responding to poor performance). (See Dominguez-Martinez *et al.* for details.)

that the net present value of the firm is maximized by  $e = 0$ . The managerial cost function,  $c(\cdot)$ , can be understood to capture how the manager herself judges the costs of such myopic actions. Critically, even if the manager internalizes some of the cost of myopic actions, in terms of equation (1) and the analysis that follows, it remains true that she still has incentives to choose an  $e > 0$ . Boosting short-term profits at the expense of long-term profits in this manner is destructive of firm value. This idea, which was developed by Stein (1989), illustrates another channel through which career concerns can lead to suboptimal decisions.

An important issue these models must address is the question of why the distortions from optimality cannot be avoided via some incentive contract. It is true that if complete contracting among *all* relevant players were feasible and unrestricted, career concerns would be unlikely to matter (e.g., all income risk could essentially be insured away). But in the real world, the relevant parties cannot fully commit to how they will use the information they learn. Because slavery is illegal and agreements among employers not to poach workers are often violations of antitrust law, it is impossible to insulate managers from trying to make themselves appealing to potential future employers. Other real-world frictions and restrictions likewise make it impossible to give real-world executives “no-dismissal” guarantees or promise them golden parachutes worth precisely equal to what they lose if dismissed. In short, it is impossible to contract away fully the negative incentives due to career concerns (see §4.1.1 of Hermalin and Weisbach, 2017a, for more on contractual solutions and their limitations).

As discussed below, as well as elsewhere in the literature (see, e.g., Hermalin, 2013, or Hermalin and Weisbach, 2017a), various economic institutions are likely to arise to compensate for an inability to write complete formal contracts. In particular, as the issue is a lack of commitment (in this simple model, a commitment as to second-period income or promise not to dismiss the manager would help to ameliorate the problem), various institutional arrangements could arise as a substitute means of providing the appropriate level of commitment. Certain aspects of governance, such as having management-friendly boards, can be thought of in this manner.

As discussed at length in §§3.2 and 4.3 of Hermalin and Weisbach (2017a), a sizeable empirical literature has arisen to test the ideas sketched above, which broadly fall under the rubric of career-concern models. For the most part, the associated empirical literature has focused on the positive incentives generated by career concerns.<sup>10</sup> But even limiting attention to the positive, the task confronting empirical scholars is daunting. Executive compensation changes from year to year for many reasons. Moreover, not all changes reflect career concerns; in particular, one should *not* assume that career-concern models substitute for or render invalid other models of managerial compensation, such as those involving formal (*à la* traditional agency theory) and informal (*à la* efficiency wage theory) incentive contracting. Instead, career concerns provide a complementary lens on compensation and

<sup>10</sup>Two notable exceptions are Edmans *et al.* (2017) and Gopalan *et al.* (2014), both of which find evidence in support of Stein’s (1989) managerial myopia model.

career-concern models are part of a large mosaic of models that help to explain corporate governance.

One strategy to overcome some of the inherent challenges of empirical work in this area is to focus on industries in which career concerns are measurable and likely to be important. One such industry is the money-management industry. As Berk and Green (2004) argue, in that industry, the market's assessment of a manager's ability should have a clear effect because it will drive inflows to the fund overseen by the manager. In particular, when a fund performs well, the market's assessment of its manager's ability increases and investors correspondingly invest more (presumably until the perceived abnormal return to managerial skill disappears). One can estimate the sensitivity of new inflows to fund performance. Since one can observe the fees charged by most funds, it is possible to estimate the additional fees a fund will earn from incremental returns, which in turn can be interpreted as the return for having a better managerial reputation. Through this process, one can, therefore, estimate the indirect, market-based, managerial incentives in the money-management industry. Moreover, these incentives are undoubtedly economically important given investors' responsiveness to their assessments of managerial ability and the close connection between size of fund and managerial compensation.

A potential downside to that estimation strategy is that it is often impossible to know what fraction of a fund's fees goes to its managers as compensation. This issue is especially true of funds managed by large organizations, such as large mutual funds. In contrast, private equity and hedge funds tend to be relatively small organizations, in which the general partners both manage the funds and receive most of the fees. These firms are, for this reason, ideal for measuring the effect of market-based incentives.

Chung *et al.* (2012) and Lim *et al.* (2016) perform this exercise for private equity funds and hedge funds, respectively. Specifically, the authors estimate the sensitivity of inflows to performance and calculate the effect on the present value of the expected fees. For both types of funds, direct-incentive compensation is large: 20% profit-sharing plus the additional incentives that derive from management ownership. These two papers use this direct incentive compensation and also calculate the "indirect Jensen-Murphy  $b$ ," the amount managers receive in future income for each additional dollar returned to investors, not including the return on the managers' own investments in the fund (for details on the Jensen-Murphy  $b$ , see Jensen and Murphy, 1990).

Table 1 shows that indirect incentives are large even relative to the large direct incentive compensation in this industry; moreover, they also represent a substantial fraction of money returned to investors. For a typical private equity fund, Chung *et al.* (2012) find that indirect incentives are nearly as large as direct incentives, while Lim *et al.* (2016) find that they are about three times as large for a typical hedge fund.

Relevant to the issues in this article, both papers find that the cross-sectional and time-series patterns of indirect incentives are consistent with notions of managerial assessment, as discussed above. For instance, in line with the theory's prediction that assessment effects diminish over a career, indirect incentives are much



**Table 1** The ratio of indirect to direct incentives and the “Jensen-Murphy  $b$ ” for different types of funds

This table presents the ratio of indirect to direct incentives and the “Jensen-Murphy  $b$ ” (in italics) for different types of funds. The “Jensen-Murphy  $b$ ” in italics is the present value of the change of managers’ revenue associated with every dollar returned to the investors (other than the managers, in the case for private equity). The left-hand side reports numbers for private equity funds from Chung *et al.* (2012), and the right-hand side reports those for hedge funds from Lim *et al.* (2016). Panel A shows the averages and Panel B the numbers by fund age or fund sequence. Numbers for private equity funds are calculated using Tables 5 and 6 in Chung *et al.* (2012), where the maximum number of future funds the GP could potentially run is three. In Panel A, it is assumed that the current fund is the first in the partnership’s sequence of private equity funds. The expected fraction of future fund sizes the GP receives as compensation is assumed to be 20% in Panels A and B. Numbers for hedge funds are calculated using Tables 4–6 in Lim *et al.* (2016), with investors’ asset value-based liquidation point,  $b$ , equal to 0.685, manager’s future expected gross-of-fee risk-adjusted performance,  $\alpha$ , equal to 3%, and the fraction of an investor’s capital that he or she withdraws each period (year) for exogenous reasons,  $\delta + \lambda$ , equal to 10%. Indirect incentives estimates are from the GIR model, which gives the most conservative measure of indirect incentives among all the models used in Lim *et al.* (2016).

	Private equity		Hedge fund	
Panel A: Averages				
Indirect/Direct	0.63		3.03	
Jensen-Murphy $b$	0.29		0.66	
	Current fund sequence		Fund age (years)	
	2	3	0	3
Panel B: By sequence/Age				
Indirect/Direct	0.47	0.32	5.51	3.33
Jensen-Murphy $b$	0.26	0.23	0.67	0.62

greater for younger partnerships than for older ones for both fund types (see Table 1). The greater uncertainty about younger partnerships’ abilities implies the market should update its assessment of management’s ability more for a given level of performance. Consequently, investment inflows and expected future fee income respond more to performance for younger than for older funds. As discussed in those articles, there are other ways in which their empirical results support an assessment or learning framework for corporate governance (in addition, also see §3.2.3 of Hermalin and Weisbach, 2017a).

Outside of the money-management industry, there have been other efforts to test and evaluate career-concern models. Boschen and Smith (1995), for example, suggest that performance does not just influence immediate or relatively contemporaneous managerial pay but the effect may last longer. Consistent with the learning framework, their estimates indicate that a onetime innovation in performance raises CEO pay for the following four to five years. They conclude that “the cumulative response of pay to performance is roughly 10 times that of the contemporaneous response” (Boschen and Smith, 1995, p. 577).

Taylor (2013) estimates market-based incentives for CEOs using a structural-estimation approach. He assumes that the CEO and board can enter into an implicit or reputational contract that permits the board to partially insure the CEO against the income risk that assessment imposes on him. His estimates suggest that 50% of any increase in firm value due to favorable news about the CEO is captured by the CEO, whereas the CEO enjoys nearly complete insurance against downward risk.

The above estimates suggest that career-based incentives can be substantial. In addition, the theory of career-based incentives has predictions about intertemporal aspects of governance and executive compensation. As alluded to above, career-based incentives should decline over time for two reasons:<sup>11</sup> first, the weight given any one period's performance declines as a manager becomes a known entity; and, second, the closer a manager is to retirement, the less time he has to capitalize on his reputation. To offset these declining indirect incentives, Gibbons and Murphy (1992) predict that optimal compensation packages have greater *direct* (i.e., performance-contingent) incentive pay components for managers close to retirement than for those far from retirement. Gibbons and Murphy (1992) test this prediction using CEO compensation data drawn from a sample of large US corporations. Their estimates indicate that the sensitivity of salary and bonus to performance is, as theory predicts, significantly greater for CEOs at the end of their careers than for CEOs in the midst of them.

However, there is undoubtedly much more to the story: other factors, related to compensation, change systematically over a CEO's tenure. In particular, a successful CEO tends to gain power *vis-à-vis* his board of directors, allowing him to capture greater compensation:<sup>12</sup> a pattern consistent with what Gibbons and Murphy find. In addition, most performance-based CEO firings occur in the first few years of a CEO's tenure (Allgood and Farrell, 2003). To induce CEOs to accept the risk inherent in such "probation," boards could make implicit promises to reward them adequately later in their career should they survive (this could be via developing a reputation to do so in a repeated-games sense or by ceding power to successful CEOs along the lines of the Hermalin and Weisbach (1998) bargaining model).

### 3. Risky Business

Learning creates volatility: stock prices respond to information learned about managerial ability, as does executive compensation and employment. Further, the more

---

<sup>11</sup>A caveat: as Holmstrom (1982) shows, if the CEO's ability is not fixed, but evolves over time according to a stochastic process, then it is possible for learning to be stationary in a way that makes incentives constant (at least if the CEO is infinitely lived). Most models in the literature, however, assume the CEO's ability is fixed after it is initially drawn.

<sup>12</sup>See Hermalin and Weisbach (1998) for a model that develops this idea. Also see §6 of Hermalin and Weisbach (2017a).

uncertainty that is being resolved, the greater is the additional volatility that learning induces.

With respect to stock returns, this idea was originally introduced by Timmermann (1993) to explain the Shiller “excess volatility puzzle.” Innovations to fundamentals can lead investors to draw inferences about future growth that, in turn, can lead to stock-price volatility that seems “too high.” Recently, Pietro Veronesi and co-authors explored the idea that learning introduces volatility in a number of different settings.<sup>13</sup>

A recent paper by Pan *et al.* (2015) extended this idea to governance in the following way. Stock prices move in response to news (good or bad) about cash flows. This movement reflects a number of things, including how the market has revised its assessment of the CEO’s ability in response to such news. As uncertainty about the CEO decreases, as it should over the course of his tenure, this component of volatility should likewise decrease. In addition, firms with CEOs of more uncertain ability will, *ceteris paribus*, have larger (in absolute value) stock-price reactions to news, and consequently higher stock return volatility.

Pan, Wang, and Weisbach find empirical support for that hypothesis: stock return volatility and the absolute value of stock price reactions to earnings news decline with CEO tenure in a convex manner. Moreover, the decline is greater when there is higher *ex ante* uncertainty about CEO ability and more transparency about the firm’s prospects. Furthermore, their estimates suggest that uncertainty about CEO ability contributes substantially to return volatility: idiosyncratic return volatility declines by 14% and total return volatility declines by 10% over the course of a CEO’s first three years in office.

Learning generates uncertainty not only for those doing the learning, but also for those being learned about. A desire to avoid such risk can distort the behavior of managers (see, e.g., the discussion of project selection in the previous section). But it can also lead those under scrutiny to demand compensation for the risk that scrutiny brings. Hence, firms can face a tradeoff between improving the quality of their assessment versus the greater compensation they must pay as a consequence of better assessment.

---

<sup>13</sup>Using a dynamic asset-pricing model, Veronesi (2000) studies the relation between stock returns and the precision of information about economic growth. Pastor and Pietro (2003) develop a model in which stock valuations depend on a learning process of average profitability. The model generates predictions that the market-to-book ratio should increase in uncertainty about average profitability, and decrease over time for each firm because of learning. Pastor *et al.* (2009) model IPO decisions in which entrepreneurs and investors learn about the average profitability from realized profits. The IPO decision for the entrepreneur is a tradeoff between benefits from diversification and those from private control. When the expected future profitability is higher than a cutoff value, it is optimal to go public. David and Veronesi (2013) construct a model where agents learn about economic and inflation regimes by observing signals such as real fundamental growth and inflation.

To explore this idea, Hermalin (2005) assumes that boards of directors vary in their diligence or independence from management. *Ceteris paribus*, a more diligent or independent board will be more willing to monitor the CEO.<sup>14</sup> Hence, a CEO facing a more diligent board faces a greater risk of dismissal than if he faces a less diligent board. All else equal, he will demand compensation for bearing that greater risk, which means his compensation will be higher than if he worked for less diligent directors. In other words, there should be a positive correlation between board independence and the level of executive compensation, a phenomenon clearly evident in time-series data.<sup>15</sup> In addition, the CEO will be tempted to mitigate the risk he faces by working harder. Although the extra effort could be beneficial for the firm, it cannot truly help the CEO in equilibrium (the Red Queen effect again); nonetheless, this extra effort reduces his utility, for which he will require offsetting compensation. Consequently, while having more diligent or independent boards might seem to be unambiguously good, in reality they can have significant costs. In particular, a likely unintended consequence of the movement by governments and exchanges to push for greater board independence is the striking increase in executive compensation over the past quarter century (see Hall and Liebman, 1998; Kaplan, 2012; Kaplan and Minton, 2012; and Peters and Wagner, 2014, for related empirical work).

Another policy that has been pushed by governments and exchanges is for better transparency and improved auditing and accounting systems within firms. Although, in light of Enron, such regulations and requirements are potentially warranted, they are not free (for many reasons).<sup>16</sup> One reason, as noted by Hermalin and Weisbach (2012), is that they increase the “scrutiny risk” to which executives are exposed. Again, executives will demand—and receive—compensation for bearing this additional risk. Thus, again, there could be an unintended consequence to this public policy: it could lead to an increase in executive compensation. In addition, greater transparency could have other pernicious effects: as Hermalin and Weisbach (2012) show, it could exacerbate the sort of value-destroying signal jamming identified by Stein (1989). The greater the monitoring (e.g., the better the disclosure), the greater the agent’s marginal benefit from concealing or distorting information and thus the greater the effort he will expend on these undesired activities. These efforts represent an additional cost to improved monitoring.

---

<sup>14</sup>It places more weight on the benefits that accrue from monitoring than its costs from engaging in that monitoring relative to a less diligent or independent board.

<sup>15</sup>See Hermalin (2005) for a detailed discussion of the associated empirical literature. Such a positive correlation is less evident in cross-sectional data. Section V of Hermalin (2005) provides details of how these results can be reconciled both with each other and with his model.

<sup>16</sup>For a discussion of some of these costs, as well as for a political-economy analysis of the principal legislative response to Enron, the Sarbanes-Oxley Act, see Romano (2005).

#### 4. Someone to Watch Over Me

Given that being assessed has substantial consequences for the CEO, it stands to reason that he would seek to affect how he is assessed. One way for him to do so is to have a say in who does the assessing. The amount of say he has depends on his bargaining power, which in turn will be a function of the outcome of his initial assessment. The higher that assessment, the greater his comparative value *vis-à-vis* any replacement, which strengthens his bargaining position with his board. He can, as indirectly modeled in Holmstrom (1982), use that power to increase his pay; but he will also wish to use to affect the way he will be assessed subsequently.

Hermalin and Weisbach (1998) develop a model along those lines. A CEO who survives an initial round of assessment is necessarily seen as superior to any potential replacement, which gives him bargaining power. Yet, he remains vulnerable to later dismissal if future assessments are less positive. He can partially protect himself against this risk by using his bargaining power to get directors who will be “friendlier” to him appointed to the board. Over time, CEOs, particularly those who have done well and, thus, have effectively gained control over their boards of directors, can leverage that control to obtain greater compensation, as well as greater freedom to pursue their own agendas. Ultimately, the market’s scrutiny of managerial abilities can go from being a means of controlling managers to one that allows managers to capture rents and possibly cause other problems for their firms.

The case of Michael Eisner, the long-time CEO of Disney, is illustrative in this regard. In the 1980s, Eisner was considered one of the best CEOs in the US, leading Disney to a 50% average annual return between 1984 and 1989. Eisner was rewarded handsomely: in each year between 1985 and 1995 he was listed by *Forbes Magazine*’s annual compensation survey as one of the 40 highest-paid CEOs between 1985 and 1995. In addition, he was able to control the process of board nominations: among those appointed to Disney’s board were the principal of his children’s elementary school, his personal attorney, and the president of Georgetown University, to which Eisner was a major contributor.<sup>17</sup> By the early 2000s, he was considered to be one of the most entrenched CEOs, attracting the ire of institutional investors. In 2005, Eisner was finally thrown out of office after a shareholder revolt and a bitter board fight.<sup>18</sup> The point of the Hermalin and Weisbach (1998) model is that the two parts of the Eisner example are fundamentally linked: by being considered a superstar, he acquired control over the board, and then used the control to further his personal agenda.

Baker and Gompers (2003), Boone *et al.* (2007), and Ryan and Wiggins (2004) each provide formal tests of the idea that successful CEOs are able to bargain for less independent boards. Boone *et al.* find that measures of CEO bargaining power and

<sup>17</sup>Reveta F. Bowers, Irwin Russell, and Leo J. O’Donovan, respectively. Sources: *Wall Street Journal*, February 20, 1998; and *The New York Times*, December 20, 1998.

<sup>18</sup>See, e.g., the *New York Times*, September 26, 2005.

tenure are negatively correlated with board independence. Similarly, Baker and Gompers find that, when firms go public, the amount of outsider representation on the board reflects their bargaining power relative to that of the CEO. Finally, Ryan and Wiggins find that CEO pay becomes less linked to firm performance as his control over the board increases (proxied by his tenure and the proportion of insiders).

In addition, a CEO's control over the board can have real costs: Coles *et al.* (2014) find evidence suggesting that boards in which the majority of directors were appointed under the current CEO ("co-opted" boards) monitor less than boards in which such directors are a minority: CEO compensation tends to be higher and performance-based turnover lower with the former than the latter. Relatedly, Pan *et al.* (2016) find that firms with co-opted boards tend to make more investments, but of lower quality, than firms without co-opted boards.

Moreover, unwinding that control is not without cost. Taylor (2010) estimates, based on their CEO-firing decisions, that boards behave as if it costs 5.9% of assets to replace a CEO. This figure amounts to a substantial quantity: \$254 million for the median-sized firm or \$1.3 billion for an average-sized firm. The reason the perceived cost to the board is so high is that it is likely to be painful for directors to fire a CEO, especially one who was responsible for their initial appointment. Taylor argues that a learning model with the assumption of a high cost of firing the CEO can explain a number of important features of the data, including the rate and timing of forced CEO turnover, the changes in profitability around forced turnovers, and the variation in profits both within and across CEOs.

A common feature of most of the literature (see, e.g., the discussion of project selection in Section 2) is that the CEO would strictly prefer that the board learn nothing about him *ceteris paribus*. If the board learns nothing, then its belief is the prior and, as discussed in the context of equation (1), that means the CEO is exposed to no risk. One can, however, conceive of situations in which the CEO might have conflicted preferences: if the board gains certain information, it will update its beliefs about him (thus exposing him to risk—a minus), but it will also be in a better position to advise him (a plus). If such advice is valuable to the firm, then the board may wish to commit to being "friendly"; that is, put in place structures (e.g., fewer independent directors, directors known to be positively biased toward the CEO, etc.) that limit its reaction if it learns negative information about the CEO. These structures provide the CEO some assurance that he can share information with the board with limited risk of repercussions. This idea that there are, thus, tradeoffs between assessing the CEO and advising him has been explored in the "friendly boards" literature, of which two key contributions are Adams and Ferreira (2007) and Harris and Raviv (2008).<sup>19</sup>

It is also worth remembering that the monitors—the board of directors—consist of individuals who also care about how they are perceived. There is evidence that

---

<sup>19</sup>For a more in-depth introduction to this literature see Adams *et al.* (2010), especially Section 2.3, as well as Adams (2017).

directors' reputations are affected by the success or failure of the firms they oversee; for example, Fich and Shivdasani (2007) find evidence that directors of firms that are alleged to have engaged in financial fraud pay a significant price in terms of getting and maintaining directorships in other firms. At a theoretical level, Hermalin and Weisbach (1998) note that a CEO dismissal could be good news about board quality because it signals that the directors are more prone to monitor. However, as Dow (2013) observes, CEO dismissal could also be a bad signal, insofar as it indicates the board is poor at assessing quality at the time of initial appointment. The fact that outsiders will draw inferences will, in turn, affect the decisions boards will make. These inferences could be positive or negative about the directors. Directors could be more vigilant about preventing financial fraud or too lenient toward the CEO to avoid looking bad for having chosen him in the first place.

## 5. Altered States

The discussion to this point has presumed that all learning is rational; that is, updating is in accord with Bayes Law. However, evidence suggests that in reality, people tend to be poor Bayesians and are prone to cognitive biases. Three such biases seem especially relevant here: the “base-rate fallacy,” the “hot-hand fallacy,” and the “fundamental-attribution bias.”<sup>20</sup> The first is a tendency to overweight the recent past and underweight base rates; that is, to overweight the signal and to underweight the prior estimate. The second, somewhat relatedly, is a tendency to forecast future success based on a recent run of success (e.g., to assume a higher probability than warranted to a basketball player's making her next shot simply because she has hit her last three).<sup>21</sup> The third is to assign more credit (or blame) to an individual and less to her circumstances than is appropriate.

The implication of the base-rate fallacy can readily be seen from equation (1), which is a weighted average of new information and the prior (the weights are  $\eta/(\eta + \rho)$  and  $\rho/(\eta + \rho)$ , respectively). A shift of weight to the first from the second will increase the CEO's incentives to affect the signal. This can be good (he works harder), bad (it distorts his investment decisions), or ugly (it reinforces his motives to pursue myopic strategies).

As discussed in §7.2.2 of Hermalin and Weisbach (2017a), the hot-hand fallacy can be understood as an extreme version of the base-rate fallacy. In the basketball context, in which the fallacy was originally observed, data reveal that a player's field goal percentage is roughly constant over time, with little to no serial correlation within a particular game. Hence, the right “model” is that each shot she takes has

<sup>20</sup>Kahneman (2011) provides an excellent summary of the literature documenting ways in which human behavior departs from rationality, including these three (and other) biases.

<sup>21</sup>The “hot-hand fallacy” was originally documented by Gilovich *et al.* (1985). A related fallacy is the false-reversion-to-mean fallacy: the erroneous prediction, for example, that someone who has had a run of heads is “due” to flip a tail on her next toss.

an independent probability of being good and that should be the probability with which one should expect her next shot to hit, even if she's enjoyed a run of success (just as one knows a fair coin will land heads with probability  $\frac{1}{2}$  no matter how many heads in a row have just been tossed). The hot-hand fallacy is, in essence, the false belief that there is something to be learned when there is, in fact, nothing to learn: as with the base-rate fallacy, sufferers underweight the prior (base rate) even though, here, the correct thing is to assign all weight to the prior.

Another way to view the hot-hand fallacy is that people are using the wrong stochastic model. Suppose, in reality, a CEO's ability is fixed over time (this, as noted earlier, is the standard assumption in the career-concerns literature). If, instead, relevant decision makers (e.g., boards of directors) mistakenly believed that ability followed a random walk (or other process with serial correlation), then those decision makers would believe that there is more to be learned from recent events than there truly is. In this regard, it is worth noting that Holmstrom (1982) considers an extension of the basic career-concerns model that encompasses abilities that follow a random walk. Although the resulting formulae are different in that version of the model than in the one with fixed ability explored earlier, the basic conclusions discussed above continue to hold—at least if abilities truly follow a random walk and this is commonly understood. What could prove interesting is to explore the deviations in beliefs, efforts, compensation, and so forth when people are subject to something like the hot-hand fallacy (falsely believe, e.g., ability follows a random walk).

The fundamental-attribution bias is more nuanced. If there is only one signal of performance, then there is no means of identifying what is attributable to the manager and what to circumstances. With more data, attribution becomes possible. For example, differences in relative performance among firms in the same industry could help identify the extent to which a given firm's success (or lack thereof) should be attributable to its CEO and the degree attributable to circumstances—say, industry-wide effects—beyond his control.<sup>22</sup> It is somewhat unclear if in practice governance properly takes into account circumstances beyond managerial control in practice. Gibbons and Murphy (1990) find evidence that relative performance is considered when boards evaluate CEOs (i.e., there is some control for circumstances), while Bertrand and Mullainathan (2001) suggest they do not.

Another issue is that if managers know they will be assessed in a relative way, then, consistent with the general point that observation can distort incentives, such assessment could lead to suboptimal behavior. For instance, Scharfstein and Stein (1990) develop a model in which relative-performance assessment yields herding in investment decisions. Making the same decision as your fellow CEOs eliminates an individual signal about you, which in turn prevents the market from updating its beliefs about your abilities. Like the project-choice models discussed above, preventing updating is a way that managers can avoid career risk.

---

<sup>22</sup>Wasserman *et al.* (2010) find evidence that who the CEO is does matter, but this effect varies considerably across industries.



There is evidence that the fundamental attribution bias applies to governance: for example, Khurana (2002a,b) has made a compelling case that there is, not infrequently, a misguided tendency to treat CEOs as rock stars (recall, too, the discussion of Eisner above). Additionally, experimental work suggests that subjects have a tendency to overweight the importance of the leader to a team's success or failure (see, in particular, Weber *et al.*, 2001).

At the same time, some caution needs to be exercised in thinking about the importance of cognitive biases. Much of the evidence from psychology and behavioral economics comes from students playing games for relatively small stakes in lab settings. Do such experiments necessarily have bearing in a corporate world in which governance mistakes have the potential to destroy billions of dollars? After all, Bayes Rule is not a secret formula and the incentives to make the right decisions are huge: shouldn't boards and others thus be strongly motivated to overcome or counteract their biases? As noted above, the empirical literature hints that the answer may nonetheless be "no," but obviously how accurately the relevant actors adjust their beliefs is an empirical question, one in need of more research.

One should also note that if such biases are truly critical, then this has implications for how we understand and assess past empirical studies. In particular, if the relevant decision makers are biased, then this raises doubts about structural estimation procedures that rely on the identifying assumption that agents rationally update beliefs using Bayes Rule.

Beyond determining the existence and magnitude of these biases in the context of governance, a number of avenues of research seem open. First, from a theoretical perspective, it is not necessarily obvious what the effect of these biases will be. For example, the base-rate fallacy causes observers to overweight current performance, which could increase a young executive's incentives relative to what they would be if observers were perfect Bayesians. On the other hand, if the past will be effectively forgotten, the long-run effects that current effort has on future compensation will be less than a model based on Bayesian updating would predict, which, in turn, reduces a young executive's incentives. Hence, one avenue of research is to determine the theoretical consequences if the relevant parties employ updating/inference rules exhibiting well-documented biases.

A second set of questions stems from the insight that if these biases are *predictable* and lead others to behave *suboptimally*, then ought there not be ways for one to make money from them? That is, why can't savvy players profit at the expense of the less savvy?

In a recent paper, Hermalin (2018) attempts to analyze some of these questions. *Inter alia*, he considers a two-period career-concerns model like that considered above, but in which one or more of the two firms bidding for the manager's services suffers from the base-rate fallacy.<sup>23</sup> As suggested above, this fallacy takes the

---

<sup>23</sup>Although, as Hermalin notes, the model can be reinterpreted in terms of the fundamental-attribution bias being the relevant cognitive bias.

form of the firms overweighting the signal and underweighting the prior. In a variant of the model without effort, he shows that the firm that suffers more from bias will lose money in expectation, except if its rival is perfectly Bayesian, in which case it breaks even. In that variant, being wholly rational does not result in making money at the expense of the less savvy. If both firms are equally poor Bayesians, they both breakeven in expectation.

When effort is reintroduced to the model, a benefit to biased updating is created for the reason sketched above: more weight on the signal raises the manager's incentive to work hard. Moreover, if each firm acts as if all players update as it does (i.e., is naïve about the existence of bias), then the more biased firm will bid more for the manager initially, thus hiring him. This firm will make money in expectation from having done so (the bidding between the firms is assumed to be a second-price sealed bid auction, somewhat akin to ascending English auction, so the winning firm captures some of the benefit of the manager's greater first-period effort). Moreover, it is possible that, in expectation, the more biased firm can earn a positive expected profit across both periods (i.e., expect to end the game having made a profit). In essence, at least in some theoretical settings, there is no reason to imagine that being a poor Bayesian will lead one to be driven from the market due to financial losses.

Part of what is going on in Hermalin (2018) is that being a bad Bayesian acts as a form of commitment to highly reward effort; this commitment is valuable and potentially profitable. Consider equation (1) again. If all firms suffered from such a severe case of the base-rate fallacy, such that  $\rho = 0$  for them, then any given CEO's return to effort is 1, matching the social return; that is, in a world of biased firms (boards of directors), the first best attains.<sup>24</sup> Admittedly, that is only a partial analysis—and certainly one can also imagine reasons why a board of bad Bayesians is bad, a point also explored in Hermalin (2018)—but it serves to illustrate one way in which cognitive biases could be a plus.

Cognitive biases can serve as valuable commitment devices in other ways.<sup>25</sup> To illustrate this possibility, another well-known learning-related bias is *confirmation bias*, the tendency to look for or overweight evidence confirming a decision and to ignore or underweight evidence contrary to it. Consider two boards, one with confirmation bias, one without. Suppose, as in Section 2, managers decide whether or not to pursue projects. The manager working for the board without bias will, for the reasons detailed in Section 2, never undertake a project. The story could be different for a manager working for a board with bias. If this board wishes to confirm the wisdom of its choice of manager, it will underweight a project failure, while

---

<sup>24</sup>This is based on the assumption of no discounting. But even allowing for discounting, the base-rate fallacy would generally seem to make incentives arising from career concerns more efficient.

<sup>25</sup>The discussion that follows is somewhat related to a strain of the leadership literature that deals with delegation. See, for example, Rotemberg and Saloner (1993, 1994, 2000), Van den Steen (2005), and Blanes i Vidal and Möller (2007).

overweighting a project success. Although choosing a project still exposes a manager to risk, the risk-adjusted upside potentially exceeds the risk-adjusted downside, which means a manager could wish to undertake projects. In sum, confirmation bias could serve as a commitment device to accept some failure, which is necessary if the manager is to be willing to pursue projects at all. As noted in §7.3 of Hermalin and Weisbach (2017a), this is a highly incomplete analysis, but it could be suggestive of possible lines of future work.

The second half of Hermalin (2018) considers how suffering from the base-rate fallacy may affect a principal's incentives to gain information about the manager. This, in turn, has bearing on the manager's effort and the risk (in that portion of the paper in terms of the manager's continued employment), which in turn bears on his compensation. Because a biased principal puts too much weight on new information, she will overinvest in collecting it. This, correspondingly, raises the manager's incentive to work harder (which benefits the principal if managerial effort is beneficial, as in Holmstrom (1982), or harms her if it is not, as in Stein (1989)), but also the manager's risk. Even if more managerial effort is beneficial, the principal could lose by overinvesting in getting new information because of (i) the cost of obtaining that information and (ii) compensating the manager for the risk; but critically, she could also gain. As Hermalin shows, there is a profit-maximizing level of bias that is positive (i.e., a somewhat biased principal will expect to earn more than a principal who is a perfect Bayesian).

## 6. Fin

The central focus of research in corporate governance has historically been on the problems of controlling managers' actions. Without minimizing the real-world importance of such control problems, we argue that such a focus is incomplete and ignores important factors affecting corporate governance. In particular, it overlooks the crucial element of career concerns: managers care about the inferences that current and future employers draw over time about their abilities from observing their performance.

Career concerns create powerful incentives for managers. Sometimes they are valuable, as Fama (1980) argues, but sometimes they can prove counterproductive, as Holmstrom (1982) and others demonstrate. In addition, because what others will ultimately infer about them is uncertain *ex ante*, managers are necessarily exposed to risk, which leads them to demand greater compensation. Further, because what is left to learn decreases over a manager's career, the corresponding uncertainty likewise decreases, which has numerous implications. Finally, how managers are assessed and by whom are endogenous decisions; decisions that are fundamentally about the design of governance structures. In short, there are a number of reasons to expect career concerns to be a major factor in corporate governance.

Of particular note is the importance of career concerns—a learning and assessment perspective more generally—for understanding the *dynamics* of corporate

governance. Such a perspective seems essential for gaining insights into phenomena such as executive selection and turnover; the evolution of executive salaries and other forms of non-contingent compensation; changes in the balance of power between boards and CEOs; and the changes in the volatility of stock returns over time.

Although by no means comprehensive—again we direct the reader to our companion piece, Hermalin and Weisbach (2017a)—our discussion of the extant literature nonetheless amply illustrates that the learning/assessment approach has already yielded considerable insights into governance. Yet this approach is not a spent force: we believe it has much to offer going forward. This is especially true for empirical work. A striking feature of the models in this literature is how well suited they are to testing: as a rule, they offer clear time-series and cross-sectional predictions. These predictions occur, in large part, because beliefs about ability are, almost by necessity, more volatile to the current performance of newcomers than of old-timers. In addition, other observables are closely correlated with factors such as the precision of the relevant prior estimate of ability or the precision of the signals available to the market. While some of these insights have already been explored empirically, many remain to be investigated.

One potential shortcoming of the current literature is that it has ignored the well-documented fact that most people are “bad Bayesians”—they fail, in systematic and predictable ways, to update their beliefs in a manner consistent with Bayes Rule. As we discussed in detail above, recognition of this fact opens a number of crucial avenues of research, both empirical and theoretical. No doubt the most critical issue in this regard is determining empirically how far real-world belief formation deviates from the Bayesian ideal in the context of governance.

In sum, learning and assessment offers a complementary approach to agency theory for the study of corporate governance; in particular, it is a means to provide a unifying framework to a number of disparate phenomena in governance. Further, it is an approach amenable to empirical analysis and, indeed, a rich empirical literature has emerged that builds implicitly and explicitly on it. Finally, it offers natural ways to incorporate behavioral insights into the study of governance.

## References

- Adams, R. B., 2017, Boards, and the directors who sit on them. In B. E. Hermalin, and M. S. Weisbach eds: *The handbook of the economics of corporate governance* (North-Holland, Amsterdam).
- Adams, R. B., and D. Ferreira, 2007, A theory of friendly boards, *Journal of Finance* 62, pp. 217–250.
- Adams, R. B., B. E. Hermalin, and M. S. Weisbach, 2010, The role of boards of directors in corporate governance: A conceptual framework and survey, *Journal of Economic Literature* 48, pp. 58–107.
- Allgood, S., and K. A. Farrell, 2003, The match between the CEO and firm, *Journal of Business* 76, pp. 317–341.

- Baker, M., and P. A. Gompers, 2003, The determinants of board structure at the initial public offering, *Journal of Law and Economics* 46, pp. 569–598.
- Becht, M., P. Bolton, and A. Röell, 2003, Corporate governance and control. In G. Constantinides, M. Harris, and R. Stulz eds: *The handbook of the economics of finance* (North-Holland, Amsterdam).
- Berk, J., and R. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, pp. 1269–1295.
- Berle, A. A., and G. C. Means, 1932, *The modern corporation and private property* (MacMillan, New York).
- Bertrand, M., and S. Mullainathan, 2001, Are CEOs rewarded for luck? The ones without principals are, *Quarterly Journal of Economics* 116, pp. 901–932.
- Blanes i Vidal, J., and M. Möller, 2007, When should leaders share information with their subordinates?, *Journal of Economics and Management Strategy* 16, pp. 251–283.
- Boone, A. L., L. C. Field, J. M. Karpoff, and C. G. Raheja, 2007, The determinants of corporate board size and composition: An empirical analysis, *Journal of Financial Economics* 85, pp. 66–101.
- Boschen, J. F., and K. J. Smith, 1995, You can pay me later: The dynamic response of executive compensation to firm performance, *Journal of Business* 68, pp. 577–608.
- Carroll, L., 1871, *Through the looking glass, and what Alice found there* (Macmillan, New York).
- Chung, J.-W., B. A. Sensoy, L. H. Stern, and M. S. Weisbach, 2012, Pay for performance from future fund flows: The case of private equity, *Review of Financial Studies* 25, pp. 3259–3304.
- Coles, J., N. Daniel, and L. Naveen, 2014, Co-opted boards: Costs, benefits, causes and consequences, *Review of Financial Studies* 27, pp. 1751–1796.
- David, A., and P. Veronesi, 2013, What ties return volatilities to price valuations and fundamentals? *Journal of Political Economy* 121, pp. 682–746.
- Dominguez-Martinez, S., O. H. Swank, and B. Visser, 2008, In defense of boards, *Journal of Economics and Management Strategy* 17, pp. 667–682.
- Dow, J., 2013, Boards, CEO entrenchment, and the cost of capital, *Journal of Financial Economics* 110, pp. 680–695.
- Edmans, A., V. Fang, and K. Lewellen, 2017, Equity vesting and investment, *Review of Financial Studies* 30, pp. 2229–2271.
- Fama, E. F., 1980, Agency problems and the theory of the firm, *Journal of Political Economy* 88, pp. 288–307.
- Fich, E. M., and A. Shivdasani, 2007, Financial fraud, director reputation, and shareholder wealth, *Journal of Financial Economics* 86, pp. 306–336.
- Gibbons, R., and K. J. Murphy, 1990, Relative performance evaluation for chief executive officers, *Industrial and Labor Relations Review* 43, pp. 30S–51S.
- Gibbons, R., and K. J. Murphy, 1992, Optimal incentive contracts in the presence of career concerns: Theory and evidence, *Journal of Political Economy* 100, pp. 468–505.
- Gilovich, T., A. Tversky, and R. Vallone, 1985, The hot hand in basketball: On the misperception of random sequences, *Cognitive Psychology* 3, pp. 295–314.
- Gopalan, R., T. T. Milbourn, F. Song, and A. V. Thakor, 2014, Duration of executive compensation, *Journal of Finance* 69, pp. 2777–2817.

- Hall, B. J., and J. B. Liebman, 1998, Are CEOs really paid like bureaucrats?, *Quarterly Journal of Economics* 113, pp. 653–691.
- Harris, M., and A. Raviv, 2008, A theory of board control and size, *Review of Financial Studies* 21, pp. 1797–1832.
- Hermalin, B. E., 1993, Managerial preferences concerning risky projects, *Journal of Law, Economics, and Organization* 9, pp. 127–135.
- Hermalin, B. E., 2005, Trends in corporate governance, *Journal of Finance* 60, pp. 2351–2384.
- Hermalin, B. E., 2013, Corporate governance. In R. S. Gibbons, and J. Roberts eds: *The handbook of organizational economics* (Princeton University Press, Princeton, NJ).
- Hermalin, B. E., 2018, Biased monitors: Corporate governance when managerial ability is mis-assessed, *Journal of the Japanese and International Economies* 47, pp. 70–80.
- Hermalin, B. E., and M. S. Weisbach, 1998, Endogenously chosen boards of directors and their monitoring of the CEO, *American Economic Review* 88, pp. 96–118.
- Hermalin, B. E., and M. S. Weisbach, 2012, Information disclosure and corporate governance, *Journal of Finance* 67, pp. 195–233.
- Hermalin, B. E., and M. S. Weisbach, 2017a, Assessing managerial ability: Implications for corporate governance. In B. E. Hermalin, and M. S. Weisbach eds: *The handbook of the economics of corporate governance* (North-Holland, Amsterdam).
- Hermalin, B. E., and M. S. Weisbach, eds, 2017b, *The handbook of the economics of corporate governance* (North-Holland, Amsterdam).
- Hirshleifer, D. A., and A. V. Thakor, 1992, Managerial conservatism, project choice, and debt, *Review of Financial Studies* 5, pp. 437–470.
- Holmstrom, B., 1982, Managerial incentive problems: A dynamic perspective. In *Essays in honor of Lars Wahlbeck* (Swedish School of Economics, Helsinki).
- Holmstrom, B., 1999, Managerial incentive problems: A dynamic perspective, *Review of Economic Studies* 66, pp. 169–182.
- Holmstrom, B., and J. Ricart i Costa, 1986, Managerial incentives and capital management, *Quarterly Journal of Economics* 101, pp. 835–860.
- Jensen, M. C., and W. H. Meckling, 1976, Theory of the firm: Managerial behavior, agency costs, and capital structure, *Journal of Financial Economics* 3, pp. 305–360.
- Jensen, M. C., and K. J. Murphy, 1990, Performance pay and top management incentives, *Journal of Political Economy* 98, pp. 225–264.
- Kahneman, D., 2011, *Thinking, fast and slow* (Farrar, Strauss, and Giroux, New York).
- Kaplan, S. N., 2012, Executive compensation and corporate governance in the U.S.: Perceptions, facts and challenges, Unpublished Working Paper, Booth School of Business, University of Chicago).
- Kaplan, S. N., and B. A. Minton, 2012, How has CEO turnover changed?, *International Review of Finance* 12, pp. 57–87.
- Khurana, R., 2002a, The curse of the superstar CEO, *Harvard Business Review* 80, pp. 60–65.
- Khurana, R., 2002b, Searching for a corporate savior: The irrational quest for charismatic CEOs (Princeton University Press, Princeton, NJ).
- Lim, J., B. A. Sensoy, and M. S. Weisbach, 2016, Indirect incentives of hedge fund managers, *Journal of Finance* 71, pp. 1623–1666.
- Milbourn, T. T., R. L. Shockley, and A. V. Thakor, 2001, Managerial career concerns and investments in information, *RAND Journal of Economics* 32, pp. 334–351.

- Pan, Y., T. Y. Wang, and M. S. Weisbach, 2015, Learning about CEO ability and stock return volatility, *Review of Financial Studies* 28, pp. 289–328.
- Pan, Y., T. Y. Wang, and M. S. Weisbach, 2016, CEO investment cycles, *Review of Financial Studies* 29, pp. 871–918.
- Pastor, L., and V. Pietro, 2003, Stock valuation and learning about profitability, *Journal of Finance* 58, pp. 1749–1790.
- Pastor, L., L. A. Taylor, and P. Veronesi, 2009, Entrepreneurial learning, the IPO decision, and the post-IPO drop in firm profitability, *Review of Financial Studies* 22, pp. 3005–3046.
- Peters, F. S., and A. F. Wagner, 2014, The executive turnover risk premium, *Journal of Finance* 69, pp. 1529–1563.
- Romano, R., 2005, The Sarbanes-Oxley act and the making of quack corporate governance, *Yale Law Journal* 114, pp. 1521–1611.
- Rotemberg, J. J., and G. Saloner, 1993, Leadership style and incentives, *Management Science* 39, pp. 1299–1318.
- Rotemberg, J. J., and G. Saloner, 1994, Benefits of a narrow business strategy, *American Economic Review* 84, pp. 1330–1349.
- Rotemberg, J. J., and G. Saloner, 2000, Visionaries, managers, and strategic direction, *RAND Journal of Economics* 31, pp. 693–716.
- Ryan, H. E. Jr, and R. A. Wiggins III, 2004, Who is in whose pocket? Director compensation, board independence, and barriers to effective monitoring, *Journal of Financial Economics* 73, 497–524.
- Scharfstein, D. S., and J. C. Stein, 1990, Herd behavior and investment, *American Economic Review* 80, pp. 465–479.
- Smith, A., 1776, *An inquiry into the nature and causes of the wealth of nations* (W. Strahan and T. Cadeli, London).
- Stein, J. C., 1989, Efficient capital markets, inefficient firms: A model of myopic corporate behavior, *Quarterly Journal of Economics* 104, pp. 655–669.
- Taylor, L. A., 2010, Why are CEOs rarely fired? Evidence from structural estimation, *Journal of Finance* 65, pp. 2051–2087.
- Taylor, L. A., 2013, CEO wage dynamics: Estimates from a learning model, *Journal of Financial Economics* 108, pp. 79–98.
- Timmermann, A. G., 1993, How learning in financial markets generates excess volatility and predictability in stock prices, *Quarterly Journal of Economics* 108, pp. 1135–1145.
- Van den Steen, E., 2005, Organizational beliefs and managerial vision, *Journal of Law, Economics, and Organization* 21, pp. 256–283.
- Veronesi, P., 2000, How does information quality affect stock returns?, *Journal of Finance* 55, pp. 807–837.
- Wasserman, N., B. Anand, and N. Nohria, 2010, When does leadership matter? A contingent opportunities view of CEO leadership. In N. Nohria, and R. Khurana eds: *Handbook of leadership theory and practice* (Harvard Business School Publishing, Boston).
- Weber, R. A., C. F. Camerer, Y. Rottenstreich, and M. Knez, 2001, The illusion of leadership: Misattribution of cause in coordination games, *Organization Science* 12, pp. 582–598.
- Williamson, O. E., 1963, Managerial discretion and business behavior, *American Economic Review* 53, pp. 1032–1057.