

Investor Adverse Selection and Corporate Governance

Paul Connell*

September 14, 2026

Abstract

If shareholders select into firms for idiosyncratic reasons, such as based on their enthusiasm for the CEO, could the shareholder base itself constitute a source of corporate governance distortion? By embedding the Holmstrom-Milgrom contracting model into a market where investors receive noisy signals about CEO skill, I show that in expectation overly optimistic investors self-select into a firm's shareholder base leading to approval of excessively steep incentive pay packages. This paper presents empirical evidence in support of that theoretical mechanism. First, steeper incentive pay correlates with the types of investors most likely to buy into a company based on internal assessments and vote their shares. Second, a structural estimation based on compensation and performance data from 2000–2025 recovers a mean investor belief error that is positive and statistically significant in a specification that permits either sign or a null result. The implied losses average \$1.16 million per firm per year, roughly sixteen percent of aggregate CEO pay. Third, say-on-pay approval, a data source external to the estimation, peaks almost exactly at the model's zero-misperception point and declines as contracts move toward either extreme, consistent with increasingly vociferous objection from less-selected members of the shareholder base as the pivotal shareholder's error grows. An asset pricing test shows that optimistically held firms suffer no return penalty, so the distortion is not a marker for larger hidden governance damages. Counterfactual policy simulations show that while a cap on incentive pay may be modestly welfare-enhancing, improving the quality of investor information is a more promising margin for intervention.

1 Introduction

In 2018, Tesla Inc. granted Elon Musk, its founder and CEO, a compensation package consisting of stock options that would vest upon achieving ambitious market capitalization and operational milestones. Musk ultimately met these aggressive goals, transforming Tesla from a \$59 billion company to one briefly worth \$1.5 trillion.¹ At recent

*University of Wisconsin Law School. Email: paul.connell@wisc.edu. I thank W. Bentley MacLeod for detailed comments on an earlier draft and participants at the 2026 American Law and Economics Association Annual Meeting and the 2026 AALS New Voices in Business Law Section for helpful discussion. During the preparation of this work, I used Claude Fable 5 for programming and editorial assistance, including code cleanup, robustness extensions and the revision of parts of the manuscript draft. I reviewed and verified the generated code and substantially revised all generated text; the analysis, arguments and errors are my own. Support for this research was provided by the Office of the Vice Chancellor for Research at the University of Wisconsin-Madison with funding from the Wisconsin Alumni Research Foundation.

¹See “Tesla Gives Musk Some Stock” Money Stuff by Matt Levine. Available at: www.bloomberg.com; published August 4, 2025, accessed September 11, 2025.

valuations, this “moonshot” package was worth approximately \$87 billion,² potentially representing an unprecedented level of executive compensation, on the one hand, and, on the other, an exemplar of efficient contracting given Tesla’s trillion-dollar value creation under Musk’s leadership.

Reacting in part to the gross magnitude of this compensation package, some shareholders sued to rescind the compensation package. In *Tornetta v. Musk*,³ the Delaware Court of Chancery invalidated it, finding that Tesla’s board lacked independence from Musk and that shareholders who approved the compensation were inadequately informed about director conflicts. Tesla’s subsequent attempts to remedy the situation—including reincorporating in Texas and securing shareholder reapproval of the initial package—initially failed to resurrect the options, with the Delaware court ruling in *Tornetta II* that retroactive ratification cannot cure the original procedural defects.⁴ This litigation culminated with *In re Tesla, Inc. Derivative Litigation*,⁵ in which the Supreme Court of the State of Delaware reversed the Chancery court’s rescission of the 2018 compensation plan on the grounds that this was an ineligible remedy and awarded the plaintiff nominal damages.

While the *Tornetta* litigation was ongoing, Tesla’s board approved an “interim” compensation package worth about \$30 billion⁶ and proposed a new pay package potentially worth up to \$1 trillion over the next ten years.⁷ Tesla’s shareholders subsequently approved this package despite opposition from several proxy advisors and certain large institutional investors.⁸

The Tesla compensation saga highlights a core tension in corporate governance: to what extent can shareholder supervision meaningfully check a powerful CEO if the shareholder base self-selected into the firm based on their belief in that CEO’s exceptional skill? In a situation in which Tesla shareholders have nothing to gain from granting additional compensation to their CEO—no level of option grants today will change Tesla’s returns over the past 7 years—they have overwhelmingly approved every compensation proposal put before them. Nevertheless, the *MFW* process that the *Tornetta* decision invoked, the subsequent amendments the Delaware legislature made to Section 144 of the Delaware General Corporation Law (DGCL),⁹ and various recent academic proposals¹⁰ all contemplate some level of non-controlling shareholder approval as a cleansing mechanism for similar transactions.

This paper contributes to the longstanding debate on CEO compensation by offering a framework that reconciles two seemingly contradictory strands of the literature. On one side, a large body of work emphasizes efficient contracting, showing that compensation levels largely track observable firm characteristics and that cross-firm patterns are broadly consistent with optimal incentive provision.¹¹ On the other side, research rooted in theories of manage-

²Id.

³*Tornetta v. Musk*, 310 A.3d 430, 445 (Del. Ch. 2024).

⁴See “Tesla Gives Musk \$30 Billion Award, Saying ‘Deal Is a Deal’” Bloomberg News, www.bloomberg.com; accessed September 11, 2025.

⁵Slip opinion dated December 19, 2025.

⁶Id.

⁷See “Tesla Offers Elon Musk an Unprecedented \$1 Trillion Pay Package”, Bloomberg News www.bloomberg.com; accessed September 11, 2025.

⁸See “Tesla Shareholders Approve Elon Musk’s \$1 Trillion Pay Package”, Wall Street Journal www.wsj.com; accessed November 17, 2025.

⁹See “What Newly Amended DGCL §144 Says (and Does Not Say) about Controlling Stockholder Transactions”, Marcel Kahan. Available at: www.corpgov.law.harvard.edu/; posted June 5, 2025, accessed September 11, 2025.

¹⁰See, for example, Goshen et al. (2024).

¹¹See, e.g., Cheng et al. (2015) who find that executive compensation packages offer appropriate risk premia for more companies with higher

rial power and rent extraction highlights that CEOs frequently receive pay unrelated to performance, often in settings marked by weak governance or informational opacity.¹² This paper shows that both patterns can emerge from the same underlying economic environment: depending on the noisiness of information available to investors, shareholder oversight can either produce well-aligned incentives or systematically distorted ones. Indeed, my empirical results suggest that well-aligned managerial incentivization is not a sufficient condition for superlative financial performance. As Section 6.2 reports, I find that firms whose shareholder bases most overvalue their CEOs suffer no measurable return penalty. In other words, a shareholder base’s overestimation of a CEO’s ability does not appear to be a “canary in the coal mine” for deeper value-destroying governance issues; a pattern perhaps exemplified by the Tesla saga.

Specifically, this paper makes three contributions. First, it develops a formal model in which adverse selection in the shareholder base leads to endogenous overvaluation of CEO ability. Investors who receive pessimistic signals do not buy into the stock, leaving behind a shareholder base composed disproportionately of optimists. These remaining shareholders rationally support compensation schemes that are excessively incentive-heavy, inducing inefficiently high effort and reducing total profits. This mechanism challenges the legal assumption—reflected among other places in Delaware’s recent revisions to DGCL Section 144—that shareholder approval reliably functions as an independent cleansing device. When the shareholder base is itself selected on optimistic beliefs, procedural protections built around shareholder consent may fail precisely when they are most needed.

Second, this paper compiles a variety of empirical data to assess the plausibility of this theoretical claim. Descriptive statistics suggest that steeper incentive pay correlates with higher levels of active institutional ownership in a company. Intuitively, these investors are the most likely to buy into a company on the basis on individual assessment of a CEO’s skill (unlike indexers who purchase a cross-section of the entire market) while simultaneously being likely to vote their shares (unlike retail investors), completing the loop required for the posited mechanism. Furthermore, using compensation and performance data from 2000–2025, the paper structurally estimates the magnitude of this adverse selection in practice. The estimated mean of investor belief errors is positive and statistically significant, and the implied losses average \$1.16 million per firm per year (median \$67,000, in 2024 dollars) or about sixteen percent of aggregate CEO compensation in the sample. In addition, I construct a dataset of 19,456 say-on-pay votes, which is a data source external to the structural estimation. The say-on-pay approval percentage, though, peaks almost exactly at the point where the structural model identifies zero misperception of CEO skill and then declines in either direction as the voted-on contracts move toward either extreme of greater error.

Third, this paper lays some initial groundwork for future inquiry. This paper’s findings of shareholder-induced CEO incentive pay distortion raise a natural question, which is whether this dynamic merits a policy response and, if so, what options are available. While the structural estimation’s aggregate recovered effect is considerable in terms

volatility, and Terviö (2008) and Gabaix and Landier (2008) who link large payouts to CEO sorting on marginal productivity that is multiplicative with firm size.

¹²See, e.g., Bebchuk et al. (2011b). See also Core et al. (1999), who find a correlation between higher CEO pay, worse financial performance and larger agency problems; Bertrand and Mullainathan (2001), who find that luck, defined as observable shocks to performance beyond a CEO’s control—such as exchange rates or oil prices—is a larger driver of CEO pay than observable changes in firm performance; and Malmendier and Tate (2009), who find that, after being so identified, “superstar CEOs” tend to under-perform.

of CEO pay, it is modest relative to firm value at roughly 0.009% of mean market capitalization ($2.9 \times 10^{-3}\%$ at the median). CEO pay may thus not drive overall returns, but a logical concern is that shareholders who are overly optimistic about CEO ability might slacken the reins in areas of oversight other than compensation, leading to further value extraction or practices that are deleterious to firm value. An asset-pricing test, however, rejects this “canary in the coal mine” hypothesis. Optimistically held firms earn ordinary risk-adjusted returns. On the other hand, a counterfactual policy experiment on the estimated model shows that a cap on incentive-pay intensity delivers modest gains, increasing firm value by an average of \$0.3 million per year per firm (or about 30% of the estimated average loss from investor misperception which is roughly \$0.46 billion per year in aggregate); though, this is admittedly an exercise in partial equilibrium. Taken together, these results indicate that interventions conditioned on observable contractual features have a low ceiling. However, interventions targeting the information environment itself remain an open question to which this paper’s measure of belief error is a portable tool for further investigation.

This paper is organized as follows. Section 2 situates the project within the existing literature. Section 3 develops the theoretical model, including the non-invertibility results that establish what market prices can and cannot reveal about CEO skill. Section 4 establishes reduced-form stylized facts linking CEO compensation and ownership structure. Section 5 implements the structural estimation of investor misperception and presents the parameter estimates and economic cost accounting. Section 6 places those estimates in their empirical context. Say-on-pay voting evidence independently locates the model’s alignment point and an asset-pricing test examines whether misperception of CEO skill correlates with other value-destroying governance practices. Section 7 develops some counterfactual policy experiments—a cap on incentive pay and an improvement in the signal quality of investor information—and the governance implications of these findings. Section 8 concludes.

2 Literature Review

This paper speaks to three strands of scholarship: (i) contract theory and the economics of incentive provision; (ii) empirical and structural work on executive compensation; and (iii) corporate governance debates regarding efficient contracting versus managerial rent extraction.

Contract theory and the principal–agent problem

The modeling framework builds on the canonical Holmström–Milgrom continuous-time principal–agent environment (Holmstrom & Milgrom, 1987), which has shaped modern understanding of optimal incentive design under moral hazard. As surveyed in Edmans and Gabaix (2016), this literature typically assumes that principals directly observe the distribution of agent types or skill parameters, even if effort is hidden. My contribution is to extend this class of models by introducing a second layer of informational friction: investors themselves receive noisy private signals about CEO ability. This generates heterogeneity in the beliefs of potential shareholders and, crucially, a form of adverse selection in who chooses to hold the firm’s securities. Embedding this friction in the Holmström–Milgrom (HM) environment

shows how standard incentive contracts can become systematically distorted even when all actors behave rationally given their information.

Empirical and structural approaches to executive compensation

A vast empirical literature studies pay-for-performance, with Edmans et al. (2017) providing a comprehensive overview. Closest in spirit to this paper is Cheng et al. (2015), who use the HM model as a benchmark to argue that cross-sectional variation in CEO pay is largely shaped by differences in firm risk. They test this proposition using compensation and returns data and find empirical support for risk-based differences. My analysis departs from this work by shifting the focus from exogenous firm risk to investor misperception of CEO ability and by employing a structural estimation approach to characterize the magnitude of these misperceptions' impact on firm performance.

Structural estimation has become increasingly prominent in the CEO-pay literature. Terviö (2008) estimates an assignment model in which CEOs and firms match based on ability and firm characteristics, concluding that much of the observed dispersion in compensation reflects heterogeneity in firms rather than managers—a finding consistent with Gabaix and Landier (2008). My model complements this work by providing a mechanism for residual variation in both pay and realized performance that arises from informational noise rather than firm characteristics, managerial skill or market structure.

Gayle and Miller (2009) and Gayle et al. (2015) similarly employ structural techniques to study moral hazard, human-capital accumulation and promotion dynamics. They show, for example, that larger firms optimally offer steeper incentive contracts and that the risk premia associated with weaker performance signals can explain a substantial share of the firm-size gradient in compensation. My contribution is complementary: rather than focusing on how signal quality varies across firms, I examine how heterogeneity in investor information can itself generate distortions in optimal contracting.

Corporate governance, institutional monitoring and rent extraction

A third strand of research debates whether observed CEO compensation reflects efficient contracting or managerial power and rent extraction. While the papers surveyed above generally emphasize that compensation levels appear broadly consistent with optimal incentive provision given economic fundamentals, a competing line of research argues that CEOs often capture rents, particularly in settings involving weak governance or informational opacity.¹³

The work in this vein closest to my project is the literature on institutional monitoring. Hartzell and Starks (2003) argue that institutional investors, by virtue of larger stakes, have stronger monitoring incentives and therefore push for both lower levels of CEO pay and higher pay-performance sensitivity. In contrast, Fahlenbrach (2009) finds that pay-performance sensitivity is negatively correlated with institutional-ownership concentration, suggesting that monitoring and incentive pay can act as substitutes. My analysis contributes to this debate by distinguishing among heterogeneous investors and identifying how imperfect information can impair institutional monitoring itself. When the shareholders

¹³See, e.g., Bebchuk and Fried (2006) and Bebchuk et al. (2011b). See also Core et al. (1999), who find a correlation between higher CEO pay, worse financial performance and larger agency problems; Bertrand and Mullainathan (2001), who find that luck, defined as observable shocks to performance beyond a CEO's control—such as exchange rates or oil prices—is a larger driver of CEO pay than observable changes in firm performance; and Malmendier and Tate (2009), who find that, after being so identified, “superstar CEOs” tend to under-perform.

who choose to hold a firm’s stock are systematically those with overly optimistic beliefs about the CEO, governance outcomes can mimic weak monitoring even if the underlying institutions are sound.

3 Model of Investor Adverse Selection

The Holmstrom-Milgrom model comes from Holmstrom and Milgrom (1987). The original model conceives of two actors: a principal who owns a firm and an agent whose effort contributes value to the firm. The principal is risk-neutral,¹⁴ but the agent is risk-averse.¹⁵ Holmstrom and Milgrom (1987) derive an optimal contract form for the principal to offer the agent in order to maximally align their incentives. This contract, c , takes the basic form of:

$$c = \phi + \theta V$$

where ϕ is a fixed dollar amount of base pay and θ is a variable pay component tied to the value of the firm, V . This solution’s realism and tractability have led to its wide adoption in the economic literature on executive compensation.¹⁶

This paper considers an HM model of CEO compensation embedded in a market in which investors operate with noisy information. Before they make their investment decisions, investors receive noisy signals about the CEO’s productivity parameter. The result of this imperfect information is that investors adversely select into the shareholder base of the firm—a “winner’s curse” result. Since the shareholder base overvalues the CEO’s productivity, they offer the CEO a contract with too much incentive pay. Over-provision of incentive pay encourages the CEO to provide more than the optimal level of effort. This increases the firm’s expected output but not enough to offset the expected increase CEO compensation, leading to a decline in profits for the firm.

3.1 Timeline

The sequence of events in this model is as follows:

- **Step 1 - Signals and valuation:** Every investor receives a noisy signal of the CEO’s productivity. Based on this signal, investors form their valuation of company shares.
- **Step 2 - Market clearing:** Investors trade shares based on their private valuations.
- **Step 3 - Contract formation:** The company’s shareholders negotiate a compensation contract with the CEO.¹⁷

¹⁴This could be, for instance, because they are a well diversified investor so they just care about maximizing the expected value of this asset and are less sensitive to variance in outcome.

¹⁵The motivation for this is that compensation from the firm might be the major source of the agent’s income, so they are more sensitive to variance in outcomes.

¹⁶See Edmans et al. (2017) for a survey of the recent literature. They note “the HM [Holmstrom-Milgrom] model has proven extremely influential due to its tractability.”

¹⁷Section 3.2.4 shows that, as investors’ preferences are single-peaked, this contract will take the form of an HM contract computed at the median shareholder’s ($\bar{i}$ ’s) valuation of the CEO’s productivity parameter ($\bar{b}_{\bar{i}}$).

- **Step 4 - Productivity realization:** Given this contract, CEO chooses their level of effort and firms realize profits/losses.

3.2 Model

3.2.1 Principal-agent problem

There is a firm whose value, V , is determined by the level of its CEO's effort, $a \in [0, \infty)$, and $\epsilon^{FIRM} \sim \mathcal{N}(0, \sigma_{FIRM}^2)$ a mean-zero error term:

$$V(a, \epsilon) = Z + ba + \epsilon^{FIRM} \quad (1)$$

where b is the marginal productivity of the CEO's level of effort and Z is the firm's size. The action, a , is taken first, then the uncertainty, ϵ^{FIRM} , is realized. None of $\{a, b, \epsilon^{FIRM}\}$ are observable to the principals (investors), though a and b are known by the agent (CEO) and the principal receives a noisy signal as to b (see Section 3.2.2).

The CEO is risk-averse with a utility function of the following form:

$$u(x) = -e^{-\eta x} \quad (2)$$

where η is a coefficient of absolute risk aversion. Exponential utility is a key assumption of the Holmstrom and Milgrom (1987) framework as it allows the removal of wealth effects. Among other things, this means that the CEO will not “satisfice” once they've achieved a certain wealth level and cease to provide additional effort.¹⁸

In addition, the agent faces an increasing cost of effort. For simplicity and to follow the exposition in Edmans and Gabaix (2016), this paper assumes a quadratic form:

$$\kappa(a) = \frac{1}{2}ga^2 \quad (3)$$

The agent's disutility from effort is assumed to be “pecuniary,”¹⁹ so we can pass (3) directly into the agent's utility function.

Given this set up, Holmstrom and Milgrom (1987) show that the optimal contract will take the form:

$$c = \phi + \theta V(a, \epsilon)$$

where ϕ is a fixed dollar amount of base pay and θ is a variable pay component tied to the value of the firm, V .²⁰

¹⁸See, for example Appendix B of Edmans and Gabaix (2016).

¹⁹Formally, the felicity function $v(c) = c$.

²⁰A derivation is provided in Appendix A.

3.2.2 Noisy investor signals

Assume there are $N \in \mathbb{N}$ investors indexed by $i \in \{1, \dots, N\}$. Each investor receives a noisy signal:

$$\tilde{b}_i = b + \epsilon_i^{INV} \quad (4)$$

Where b is the true productivity of the CEO's effort and $\epsilon^{INV} \sim \mathcal{N}(0, \sigma_{INV}^2)$ is a mean-zero error term.

Based on the receipt of these signals, each investor derives what they believe would be the optimal contract for the CEO given their perceived productivity parameter, $\tilde{b}_i$. Based on this, they forecast the CEO's provided effort level, the resulting firm value and profits. That is, each investor forms the expectation

$$\mathbb{E}_i \left[\frac{\tilde{V}_i(\tilde{a}_i, \epsilon_{FIRM}) - (\tilde{\phi}_i + \tilde{\theta}_i \tilde{V}_i(\tilde{a}_i, \epsilon_{FIRM}))}{s} \right] \quad (5)$$

where $\tilde{\phi}_i$ and $\tilde{\theta}_i$ are the parameters for the optimal HM contract given productivity parameter $\tilde{b}_i$, $\tilde{a}_i$ is the level of agent effort the investor expects this contract to induce, $\tilde{V}_i$ is the resulting expected firm value and $\tilde{\pi}_i = \tilde{V}_i(\tilde{a}_i, \epsilon_{FIRM}) - (\tilde{\phi}_i + \tilde{\theta}_i \tilde{V}_i(\tilde{a}_i, \epsilon_{FIRM}))$ are the expected profits (firm value minus compensation costs). The firm is assumed to have $s \in \mathbb{N}$ shares and profits are distributed evenly across shares. Hence, investor i is willing to pay up to $\frac{\mathbb{E}_i[\tilde{\pi}_i]}{s}$ for a share of the company.

3.2.3 Market-clearing

To identify the market equilibrium for the company's shares, I make two simplifying assumptions.

First, I assume that each investor can own a maximum of one share. Practically, this allows me to avoid modeling investor-specific budget constraints. If investors could buy more than one share, the investor with the highest valuation would saturate their budget constraint, then the investor with the second-highest valuation would do so, and so on, until no shares remained. Concentrating the shareholder base in fewer, more extremely optimistic hands would only heighten the adverse selection the model is designed to highlight, so this assumption is, if anything, conservative.

Second, I assume that investors trade on their own signal-based valuations and do not revise their beliefs about CEO skill upon observing the market price or their counterparties' willingness to trade. This is the standard primitive of the difference-of-opinion literature,²¹ and it is less restrictive here than it first appears. Proposition 1 below establishes that the equilibrium price is consistent with any value of the CEO's true skill, so the assumption at hand does not foreclose an opportunity that some actor may have had to profit by updating their valuation belief.

Given these assumptions, the market for shares is an assignment game in the sense of Shapley and Shubik (1971). More to the point, it is the special, homogeneous case of the assignment game because the shares are identical. The surplus any investor derives does not depend on which specific share they hold. The s shares are supplied inelastically by initial owners whose reservation values lie below investors' valuations, and investors' demands are unit (one or

²¹See, e.g., Miller (1977).

none). Hence, an equilibrium is an allocation of the s shares and a price $p^\star$ at which no blocking pair remains—i.e., no excluded investor and current holder can both gain from making a trade. For this market, the core delivers two conclusions.²² First, the allocation is unique up to ties, which occur with probability zero under continuously distributed signals. In equilibrium, the shares will be held by the investors who have the s highest valuations $\frac{\mathbb{E}_i[\tilde{\pi}_i]}{s}$. Second, as the shares are identical, they must trade at a common price in equilibrium. If two shares were to trade at different prices, the buyer paying more and the seller receiving less would themselves constitute a blocking pair. The set of prices consistent with the homogeneous assignment game’s core is thus the interval between the valuation of the marginal holder (the s^{th} highest) and that of the best excluded investor (the $(s + 1)^{th}$ highest). Hence, the market-clearing price satisfies

$$p^\star \in \left[\frac{\mathbb{E}_{(s+1)^{th}}[\tilde{\pi}_{(s+1)^{th}}]}{s}, \frac{\mathbb{E}_{s^{th}}[\tilde{\pi}_{s^{th}}]}{s} \right] \quad (6)$$

and each of the investors with the s highest valuations buys a share and expects, under their own beliefs, a surplus of

$$\frac{\mathbb{E}_i[\tilde{\pi}_i]}{s} - p^\star \geq 0.$$

Where analysis in this paper requires a point value for $p^\star$, such as in the price-inference results of Section 3.2.7, I take $p^\star$ to be the marginal holder’s valuation at the upper endpoint of the interval.

3.2.4 Contract-setting

Once the shareholders have been determined, they have to set their contract with the CEO to manage the firm. As the shareholders differ only in their assessments of the CEO’s marginal productivity (their $\tilde{b}_i$), their preferences over contract slopes are single-peaked, with each shareholder’s peak at the slope optimal under their own $\tilde{b}_i$. Figure 1 provides a graphical depiction and a proof is provided in Appendix B.²³ Therefore, the median voter theorem applies and the contract that would win a shareholder vote is the one proposed by the median—as ranked by $\tilde{b}_i$ —shareholder. I will denote this median shareholder $\bar{i}$. To formalize this intuition about the contract-setting process, assume further:

Assumption 1 (Contract formation). *Shareholders vote over the incentive slope, θ . Given the implemented slope, base pay, ϕ , is determined in bilateral negotiation between the firm and the CEO, holding the CEO to their reservation utility, ω , at the CEO’s true productivity $b^\star$. Neither ω nor $b^\star$ is observed by investors, who only observe the signed contract (ϕ, θ) .*

Assumption 1 has a few functional implications. When voting, each shareholder forecasts both the CEO’s effort

²²Shapley and Shubik (1971) develop the general two-sided assignment game and show its core coincides with the set of competitive price vectors; the homogeneous case is the “horse market” of Böhm-Bawerk with which they open, where core prices are bounded by the marginal pairs.

²³Each investor takes their own $\tilde{b}_i$ as the truth and therefore evaluates contracts proposed under others’ beliefs $\tilde{b}_{-i}$ as the off-optimum points along the green curve in Figure 1.

and the negotiated base pay under their own belief $\tilde{b}_i$. This delivers the cross-evaluation profit used in the single-peakedness proof.²⁴ Realized profit is evaluated at the negotiated $\phi(\theta, b^*)$, so the CEO's participation is individually rational by construction. Finally, investors' observation of the signed contract does not reveal b^* , because ω and b^* enter base pay jointly. So, with one observation and two unknowns the investors cannot back out either variable.²⁵

3.2.5 Productivity realization

Assuming that the agent knows their true productivity, b^* , they will choose activity level:

$$a^* = \frac{\tilde{\theta}_i b^*}{g} \quad (7)$$

Proposition 2 (proved in Appendix B) establishes that, under the conditions stated therein (in particular, noisy signals and excess demand for shares), the median shareholder's signal is biased upward. That is, $\mathbb{E}[\tilde{b}_i] > b^*$ and $P(\tilde{b}_i > b^*) > 1/2$. Corollary 1 shows that this bias translates into excessive incentive pay, inefficiently high CEO effort and therefore reduced firm profit. Of note, the preceding statements are all with regard to expected-value. For any particular realization of investor signals, the median shareholder may under- or overestimate CEO skill. However, the systematic tendency is toward overvaluation. This asymmetry is the core mechanism depicted in Figure 1—i.e., that adverse selection pushes the shareholder base to the right of the profit-maximizing contract.

Figure 1 portrays the expected results of implementing HM contracts at various “misperceived” values, $\tilde{b}$, of b^* . For values $\tilde{b} < b^*$, the CEO's productivity is undervalued, so not enough incentive pay is offered. The CEO's effort level is therefore suboptimal and, while compensation costs decrease relative to the optimal ($\tilde{b} = b^*$) level this is more than offset by a lower firm output. Conversely, for values $\tilde{b} > b^*$, the CEO's productivity is overvalued and too much incentive pay is offered. While this results in higher levels of effort by the CEO, the resulting increase in firm value does not outstrip the increase in compensation cost, so total profits decrease.

3.2.6 Price Non-Invertibility and Belief Persistence

A natural concern with the adverse-selection mechanism described above is whether market prices themselves would correct it. If the market-clearing price were fully revealing in the sense that an investor could infer b^* from observing p^* , then each shareholder would update their belief to the CEO's true productivity and the distortion would unravel before the incentive contract is made. This subsection shows formally that absent perfect information about every aspect of the trading environment—including such details as the total number of potential traders or the exact dispersion of the noise in investors' signals (σ_{INV})—prices are not fully revealing and belief updating may not occur.

²⁴See Appendix B.

²⁵A foreseeable question is whether the CEO could unwind the distortion by simply revealing b^* . Within the described model there is no strict incentive for the CEO to do so. Base pay binds the participation constraint at their true productivity (Assumption 1), so the CEO's certainty equivalent is ω under any implemented slope. Hence the revelation would have no private value. Furthermore, an announcement would be merely “cheap talk” as this model admits no certifiable disclosure mechanism other than realized performance, which does not occur until after the contract is set.




Figure 1: Costly Misperception of CEO Productivity. This figure shows the decomposition of firm profit into firm value and CEO compensation as a function of investor perception of CEO productivity. When investors overestimate CEO skill ($\tilde{b}/b^* > 1$), they set excessive compensation that reduces firm profit despite increased firm value. All values are normalized by the profit at optimal perception. The vertical dashed line indicates correct perception where $\tilde{b} = b^*$. Parameters used to generate the graph were: $[b^* = 4, \sigma_{FIRM} = 2, g = 2, \eta = 2, \omega = 0]$.

I then discuss the relationship between these results and the broader literature on limits to arbitrage and information aggregation.

Recall that each investor i forms a per-share valuation:

$$w_i \equiv \frac{\mathbb{E}_i[\tilde{\pi}_i]}{s}$$

where $\tilde{\pi}_i$ is the expected profit computed at investor i 's perceived productivity $\tilde{b}_i = b^* + \epsilon_i^{INV}$. Since w_i depends on the model parameters $(\sigma_{FIRM}, g, \eta, Z, s)$ —parameters common to all investors²⁶—and is strictly increasing in $\tilde{b}_i$,²⁷ ranking investors by valuation is equivalent to ranking them by signal. The market-clearing price is therefore pinned down by the s -th highest signal realization among N draws from $\mathcal{N}(b^*, \sigma_{INV}^2)$ (under the selection mechanics of Section 3.2.3).

Let $\tilde{b}_{(s:N)}$ denote the s -th order statistic in descending order of N i.i.d. draws from $\mathcal{N}(b^*, \sigma_{INV}^2)$. By the monotonicity of $w(\cdot)$, observing p^* is equivalent to observing $\tilde{b}_{(s:N)} = \hat{b}$ for some $\hat{b} \in \mathbb{R}$. The density of this order statistic is:²⁸

$$f_{(s:N)}(\hat{b} \mid b^*, N) = \frac{N!}{(s-1)!(N-s)!} \left[\Phi\left(\frac{\hat{b} - b^*}{\sigma_{INV}}\right) \right]^{N-s} \left[1 - \Phi\left(\frac{\hat{b} - b^*}{\sigma_{INV}}\right) \right]^{s-1} \frac{1}{\sigma_{INV}} \phi\left(\frac{\hat{b} - b^*}{\sigma_{INV}}\right) \quad (8)$$

where Φ and ϕ denote the standard normal CDF and PDF, respectively.

This subsection proceeds with two stages of results. Proposition 1 establishes that the market-clearing price is globally non-invertible if the exact number of market participants is unknown. That is, any value of b^* assigns positive density to the observed price. Section 3.2.7 then provides a stronger local result. Lemma 1 shows that whenever any single environmental parameter is unknown, alternative scenarios exist on both sides of the truth that are statistically indistinguishable from it in likelihood-ratio terms. Appendix B shows that this result holds when the parameter in question is market thickness, signal precision, effective float or baseline firm value.

Proposition 1 (Non-invertibility of the market-clearing price). *Let $s \geq 1$ be the number of shares and $N \geq s$ the number of potential investors, and suppose investor i observes the market-clearing price p^* (equivalently, the signal realization $\tilde{b}_{(s:N)} = \hat{b}$ that determines it). Then for any observed $\hat{b} \in \mathbb{R}$, every candidate pair (b', N') with $b' \in \mathbb{R}$ and $N' \geq s$ assigns positive density to $\hat{b}$.*

Proof. Fix any $\hat{b} \in \mathbb{R}$, any $b^* \in \mathbb{R}$, and any $N \geq s$. Since $\phi(\cdot) > 0$ on $\mathbb{R}$ and $\Phi(\cdot) \in (0, 1)$ for all finite arguments, every factor of equation (8) is strictly positive, so $f_{(s:N)}(\hat{b} \mid b^*, N) > 0$. Hence the observed price is consistent with every (b^*, N) , and in particular with values of b^* arbitrarily far from the truth. $\square$

The Proposition 1 establishes that no value of CEO skill can be ruled out on the basis of the market-clearing price

²⁶For Proposition 1, these are taken as known to the inferring investor. The corollaries of Section 3.2.7 relax several of them, along with N and σ_{INV} , one at a time.

²⁷This follows from the fact that higher perceived productivity implies a steeper optimal contract, higher expected effort, and higher expected firm value net of compensation. See Lemma 3 in Appendix B.

²⁸Re-indexing from ascending order: the s -th from the top among N draws is the $(N - s + 1)$ -th from the bottom.

alone, even when all environmental parameters are known. Furthermore, the next subsection shows that this result is not economically vacuous in the sense that there exists a neighborhood of the true parameters in which alternative parameter combinations are so close in likelihood to the truth that a Bayesian investor could not even approximately identify b^* .

3.2.7 Generalized Non-Invertibility

The non-invertibility result is not special to market thickness (i.e., the relative magnitudes of N and s). The market-clearing price is a single number that reflects the s -th order statistic of the investors' signals. On the other hand, the environment that generates this price has several parameters, including the number of potential investors N , the effective share supply s , the precision of investors' signals σ_{INV} and the firm's baseline value Z . Inverting the price for b^* therefore requires the investor to be certain as to each of these parameter's exact value. Lemma 1 shows that if any such value is unknown, the price is globally uninformative about b^* (in the sense that no level of CEO skill can be ruled out) and locally uninformative in likelihood-ratio terms (alternatives on either side of the truth fit the observed price exactly as well). Full revelation is thus a knife-edge condition of perfect information.

Lemma 1 (Generic non-invertibility of the market-clearing price). *Let the market-clearing price be determined by $\hat{b} = \tilde{b}_{(s:N)}$ as above, and let ξ denote one environmental parameter, $\xi \in \{N, s, \sigma_{INV}, Z\}$, with true value ξ_1 interior to its domain (integer parameters extended to the reals via the Gamma function). Write $\ell(b^*, \xi)$ for the log-likelihood that the investor assigns to the observed price under the pair (b^*, ξ) , holding all other parameters at their true values. Then:*

- (i) **(Global)** *For every $b^* \in \mathbb{R}$ and every admissible ξ , the observation has strictly positive likelihood. (This restates Proposition 1 in the present notation.)*
- (ii) **(Local)** *If, moreover, $\partial \ell / \partial \xi \big|_{(b_1^*, \xi_1)} \neq 0$, then there exists an open interval $\mathcal{U} \ni b_1^*$, containing points on both sides of b_1^* , and a C^1 map $\xi^* : \mathcal{U} \rightarrow \mathbb{R}$ with $\xi^*(b_1^*) = \xi_1$, such that*

$$\ell(b_2^*, \xi^*(b_2^*)) = \ell(b_1^*, \xi_1) \quad \forall b_2^* \in \mathcal{U}.$$

That is, every nearby skill level is exactly as likely as the truth once the unknown parameter is adjusted, so the likelihood ratio provides no discriminating power.

Proof sketch. Global positivity follows, as in Proposition 1, because the likelihood is a product of strictly positive factors at finite arguments. For the local statement, the log-likelihood of the observed price is a smooth function of (b^*, ξ) and, by hypothesis, strictly sensitive to the environmental parameter at the truth. The Implicit Function Theorem therefore yields a level curve through the true parameter pair along which the likelihood is constant. As perceived skill moves away from the truth in either direction, the unknown parameter can be adjusted to fit the observed price with

equal likelihood. Appendix B contains the full proof, together with the explicit construction for the market-thickness parameter and including the treatment of integer-valued thickness. $\square$

For application of the Lemma to each of the environmental parameters, the only work is to verify that the derivative $\partial \ell / \partial \xi$ is nonzero at the truth. The corollaries in Appendix B display the relevant derivatives in closed form (each of which vanishes only on a measure-zero set of parameter configurations).

As a further observation, if $k \geq 2$ of the environmental parameters are unknown simultaneously, the set of $(b^*, \xi_1, \dots, \xi_k)$ configurations exactly as likely as the truth is generically a k -dimensional manifold through the true parameter vector. In other words, each additional unknown parameter value adds a dimension along which the truth can be traded off, implying an even less favorable environment for inverting true CEO skill from the market-clearing price.

3.2.8 Implications and Limitations of the Model

The analysis above is subject to a few limitations which are worth discussing before proceeding to the empirical implementation of the model. For one thing, the preceding analysis focuses exclusively on price signals. There are, however, other conduits for information such as public disclosure documents, media coverage, analyst reports and so on. Nonetheless, these signals are either public, in which case they are common to all market participants and shift the entire signal distribution (without revealing any individual's location within the distribution), or they are private, in which case they are definitionally a component of the individual's noisy signal $\tilde{b}_i$.²⁹

Another aspect of the model design meriting further discussion is the model's abstraction of the CEO compensation-setting process. While pay packages are generally negotiated between the CEO and the board of directors, I chose instead to collapse this pay-setting process into a direct shareholder choice. While this choice loses some fidelity, it fits the paper's motivating question regarding *Tornetta* (which boiled down to a shareholder vote on CEO incentive pay)³⁰ and it has reasonable theoretical and empirical grounding. Theoretically, investor selection ought to propagate through board elections as directors are elected by the same adversely-selected shareholder base. Leveraging the paper's same median-voter logic over board slates, the compensation committee should be an agent of the selected median shareholder. The model's comparative statics, then, should carry through with any board objective that is increasing in the median shareholder's perceived CEO skill. Empirically, there is evidence that boards respond to shareholder pay preferences. The say-on-pay literature (e.g., Thomas et al. (2012), Ferri and Maber (2013), Fisch et al. (2018)) documents boards cutting or restructuring pay after low advisory votes.³¹ Hence, the shareholder preference

²⁹Additionally, the structural estimation of CEO skill is not in tension with the non-invertibility results of Section 3.2.6, which concern the spot price of a company's stock. The estimation proceeds by leveraging the CEO's contract's slope to reveal the pivotal shareholder's belief and by estimating the CEO's true productivity based on decades of realized performance. I.e., it is an *ex post* reconstruction from a long panel of data that market participants do not possess in real time.

³⁰A shareholder vote on pay is also the mechanism at play in the revised DGCL §144 / MFW cleansing votes and the say-on-pay data discussed below in Section 6.1.

³¹Furthermore, as Section 6.1, below, reports, shareholder say-on-pay votes track the model's estimated misperception. Since the say-on-pay data were external to every input of the structural estimation, this alignment lends confidence to the underlying assumption that CEO compensation aligns with shareholder preferences.

that this paper’s model characterizes is plausibly the relevant metric in analyzing CEO pay even when the board is the agent involved in the actual negotiation.

Finally, this paper’s non-invertibility results complement a large theoretical and empirical literature documenting that arbitrage is limited in practice. Capital constraints, noise-trader risk, funding risk, short-sale frictions and volatility concerns all impede the ability of sophisticated traders to push prices toward fundamental values.³² In the theoretical limit, perfectly informed arbitrageurs with unlimited wealth and risk tolerance would eliminate mispricing. But this benchmark implicitly assumes the existence of a perfectly informative signal about managerial skill—precisely the signal that Proposition 1 shows cannot be extracted from the market price alone.³³ Additionally, in other classical non-revelation results (e.g., Grossman and Stiglitz (1980)), prices fail to aggregate information because exogenous noise trading obscures the informed order flow. This paper’s results require no noise traders and non-revelation arises instead because selection makes the price an order statistic, which confounds the location of the signal distribution with its thickness (N), rank (s) and scale (σ_{INV}). Adding noise traders would only strengthen the conclusion.

4 Reduced-Form Evidence on Investor Base and CEO Pay

4.1 Stylized Facts Motivating the Structural Estimation

The purpose of this section is to document several empirical relationships that align with the theoretical model developed above. These reduced-form patterns help establish that the model’s key mechanisms appear sufficiently present in the data to justify a structural estimation of the model’s parameters.

I assemble data on executive compensation from ExecuComp, stock market data from CRSP and institutional ownership data from FactSet, all accessed through WRDS. The sample period for this reduced form analysis spans 2013–2022, beginning with the first full year in which public firms were subject to comprehensive online reporting of Section 13 ownership filings.

4.1.1 Baseline Relationship: Investor Composition Correlates with CEOs’ Incentive Compensation

The theoretical model predicts that incentive pay—which is driven by the contract slope θ in the Holmström–Milgrom framework—should increase with the extent to which investor adverse selection upwardly biases the shareholder base’s belief about CEO ability. To proxy for investor selection, I use the fraction of a firm’s shares held by *active* institutional investors (i.e., non-indexing, non-retail owners). Active investors are the relevant group for two reasons: (i) retail shareholders typically do not vote and therefore play little role in setting CEO pay regardless of their

³²See, e.g., Shleifer and Vishny (1997), De Long et al. (1990) and Hong and Stein (2007).

³³A related objection is that pessimistic investors ($\epsilon_i^{INV} < 0$) might short the stock, placing downward pressure on the price. However, if short positions are permitted, there is no reason to exclude margin lending to optimistic investors, allowing them to lever up and counteract pessimistic shorting. The net directional effect is therefore ambiguous *ex ante*. More fundamentally, profitable shorting requires identifying the overpricing, which Proposition 1 shows cannot be extracted from the assumed environment. Empirically, stocks with substantial short interest—such as Tesla, the motivating example for this paper—frequently exhibit extended periods of valuations well above fundamental benchmarks. See, e.g., Sam Pierson, “Tesla: Most Shorted Stock Ever?,” S&P Global, available at www.spglobal.com.

beliefs, and (ii) passive indexers hold largely fixed portfolios across firms, making their ownership unlikely to reflect firm-specific misvaluation.

Formally, the baseline specification is:

$$\begin{aligned} \text{var_comp}_{it} = & \beta_0 + \beta_1 \text{act_io}_{it} + \beta_2 \ln(\text{mkt_cap}_{it}) \\ & + \sum_t \lambda_t \mathbf{1}\{\text{Year} = t\} + \sum_k \delta_k \mathbf{1}\{\text{Industry} = k\} + \sum_j \gamma_j \mathbf{1}\{i = j\} + \varepsilon_{it} \end{aligned} \quad (9)$$

where λ_t denotes year fixed effects, δ_k denotes industry fixed effects based on 6-digit NAICS codes, and γ_j denotes firm fixed effects. The variable “ act_io_{it} ” is the percent of the shareholder base of firm i that is held by active investors; column (2) replaces it with its one-year lag, measured the year before the compensation package is realized. The variable “ $\ln(\text{mkt_cap})_{it}$ ” is the natural logarithm of company i ’s market cap in year t .

Columns (1) and (2) in Table 1 provide the estimation results of equation (9). Column (1) uses contemporaneous active investor ownership, while column (2) uses its one-year lag. Ideally, ownership shares would be measured at the time of proxy voting, but proxy season varies across firms. The consistency of coefficient magnitudes and significance across specifications therefore provides reassurance that the variation in active ownership relevant for compensation decisions is relatively stable over the time period I care about. Both specifications show statistically significant positive associations with similar magnitudes between active ownership and the share of CEO pay that is variable.

4.1.2 Arbitrageur Presence does not Correlate with CEO Incentive Compensation

The next stylized fact concerns whether the relationship between ownership composition and variable pay might instead be driven by arbitrageurs. If arbitrage activity efficiently corrects investor misperceptions, then observed compensation variation could be explained by arbitrage-induced pricing dynamics rather than the adverse-selection channel emphasized in the model. To assess this possibility, I augment the baseline regression with empirical proxies for arbitrage activity:

$$\begin{aligned} \text{var_comp}_{it} = & \beta_0 + \beta_1 \text{act_io}_{i,t-1} + \beta_2 \ln(\text{mkt_cap}_{it}) \\ & + \beta_3 (\text{Bid-Ask Spread}_{it}) \\ & + \beta_4 \ln(\text{Yearly Avg. Trading Volume}_{it}) \\ & + \beta_5 (\text{Stock Price Volatility}_{it}) \\ & + \sum_t \lambda_t \mathbf{1}\{\text{Year} = t\} + \sum_k \delta_k \mathbf{1}\{\text{Industry} = k\} + \sum_j \gamma_j \mathbf{1}\{i = j\} + \varepsilon_{it} \end{aligned} \quad (10)$$

Column (3) of Table 1 shows that conventional indicators of arbitrage activity are not driving the baseline relationship. Bid–ask spreads are essentially uncorrelated with variable compensation, and stock price volatility and trading volume—which are typically inversely related to arbitrage engagement—enter positively. This pattern is consistent

with CEOs demanding higher variable compensation for bearing higher firm-specific risk (see Cheng et al. (2015)), rather than with arbitrage correcting mispricing.

4.1.3 CEO Tenure Correlates with More Incentive Pay

The theory also predicts that incentive pay should rise with investors' perceptions of CEO skill. CEO tenure serves as a useful proxy for perceived ability: poorly performing CEOs are less likely to remain in office, creating a survivorship channel. Accordingly, I estimate:

$$\begin{aligned} \text{var_comp}_{it} = & \beta_0 + \beta_1 \text{act_io}_{i,t-1} + \beta_2 \ln(\text{mkt_cap}_{it}) \\ & + \beta_3 \text{CEO Tenure}_{it} \\ & + \sum_t \lambda_t \mathbf{1}\{\text{Year} = t\} + \sum_k \delta_k \mathbf{1}\{\text{Industry} = k\} + \sum_j \gamma_j \mathbf{1}\{i = j\} + \varepsilon_{it} \end{aligned} \quad (11)$$

Column (4) of Table 1, which reports the results of estimating equation (11), confirms that longer-tenured CEOs receive a greater share of compensation in the form of incentives, which is consistent with the model's prediction that higher perceived ability leads to steeper optimal contracts.

4.1.4 Reverse-Causality Check: Higher CEO Incentive Pay Does not Predict Changes in Actively Managed Shareholdings

Finally, I address the concern that the causality may run in the opposite direction—i.e., that investors might prefer firms whose CEOs already have high incentive pay. To test this, I estimate:

$$\begin{aligned} \text{var_comp}_{it} = & \beta_0 + \beta_1 \text{act_io}_{i,t-1} + \beta_2 \text{act_io}_{i,t+1} + \beta_3 \ln(\text{mkt_cap}_{it}) \\ & + \sum_t \lambda_t \mathbf{1}\{\text{Year} = t\} + \sum_k \delta_k \mathbf{1}\{\text{Industry} = k\} + \sum_j \gamma_j \mathbf{1}\{i = j\} + \varepsilon_{it}. \end{aligned} \quad (12)$$

Regression (5) shows a strong positive association between past active ownership and current variable compensation, mirroring the baseline results. In contrast, future active ownership (“*act_io_{t+1}*”) is statistically insignificant and indistinguishable from zero. This pattern suggests that active investors do not appear to be selecting into firms because of high incentive compensation, reducing the likelihood of reverse causality.

Table 1 reports the full set of regression estimates.

Table 1: Reduced Form Regression Results

	<i>Dependent variable: Variable Compensation as a Percent of Total Compensation</i>				
	Baseline (1)	Baseline Lag (2)	Arbs (3)	CEO Tenure (4)	Leads and Lags (5)
Constant	0.253*** (0.040)	0.266*** (0.044)	0.152* (0.081)	0.273*** (0.044)	0.304*** (0.051)
Active Investor Percent (act_io)	0.123*** (0.032)				
ln(Market Cap)	0.045*** (0.004)	0.044*** (0.005)	0.044*** (0.005)	0.043*** (0.005)	0.039*** (0.005)
Bid-Ask Spread			-0.000 (0.000)		
ln(Yearly Avg. Trading Vol.)			0.011* (0.006)		
Stock Price Volatility			0.000*** (0.000)		
CEO Tenure (yrs.)				0.002*** (0.001)	
act_io _{t+1}					0.036 (0.036)
act_io _{t-1}		0.133*** (0.032)	0.127*** (0.032)	0.131*** (0.032)	0.103*** (0.035)
Year FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Observations	15,521	13,517	13,516	13,439	11,603
R ²	0.662	0.681	0.682	0.683	0.702
Adjusted R ²	0.612	0.628	0.629	0.630	0.647

Notes: OLS estimates on the CEO firm-year panel, 2013–2022 (ExecuComp × FactSet × CRSP). Dependent variable: variable compensation as a share of total compensation, (tdc1 – total_curr)/tdc1. Active Investor Percent is total institutional ownership minus top-3 passive ownership; continuous regressors winsorized at the 1st/99th percentiles. Lagged and lead ownership are one-year firm-level lags/leads on the CEO firm-year panel. Standard errors clustered by firm (PERMCO) in parentheses. *p<0.1; **p<0.05; ***p<0.01.

5 Structural Estimation

5.1 Overview

The following subsections provide the mathematical details of the structural model and its estimation. This subsection provides a brief, non-technical overview of the estimation strategy. The goal of the estimation is to recover two objects that are not directly observable in the data: (i) how much investors misperceive CEO ability, and (ii) how much of firms' performance volatility is due to how volatile firm fundamentals are, against which realized performance is judged. This paper's method proceeds from the assumption that modern CEO pay contracts embed a precise mathematical relationship between the level of pay-for-performance incentives and the perceived ability of the manager. In the Holmström–Milgrom framework, this relationship runs in one direction (from beliefs to contracts), and the model as described above allows me to invert it. That is, given an observed compensation contract, assuming the model above is a correct description of the world, I can infer what level of managerial ability investors must have believed in when writing that contract.

The estimation proceeds in two steps. First, for each firm, I use observed incentive contracts to back out a time series of perceived managerial skill. Second, I compare those perceived skills to the firm's actual performance over time. Firms whose compensation contracts imply high perceived skill should—if those beliefs were accurate—tend to generate reliably higher performance. When they do not, however, the gap between perceived and true skill may be interpreted as investor misperception. By aggregating these gaps across many firms and many years, the model identifies the overall distribution of investor misperception and the amount of fundamental noise in firm outcomes.

In practical terms, the structural estimation chooses the values of the three key parameters that make the entire data set “fit” the model most closely. These are: the volatility of firm-level shocks (σ_{FIRM} in the model), the dispersion of investor belief errors (σ_{INV}) and the average belief error (μ_{INV})—whose sign (positive, consistent with overvaluation) is the model's central prediction. For each candidate parameter vector ($\sigma_{FIRM}, \sigma_{INV}, \mu_{INV}$), the estimation procedure solves a simple optimization problem for each firm, selecting the level of true managerial skill that best reconciles (i) the incentives investors wrote into the contract and (ii) the firm's realized performance. I employ a likelihood function which aggregates these firm-level fits and selects the parameter values that globally rationalize the observed pattern of CEO contracts and outcomes.

5.2 Contracting Environment

The contracting environment follows the Holmström–Milgrom model. Firm value evolves as

$$\Delta V_{it} = \theta_{it} (b_i^*)^2 + \epsilon_{it}^{FIRM} - c_{it}(V_{it})$$

where $\epsilon_{it}^{FIRM} \sim \mathcal{N}(0, \sigma_{FIRM}^2)$ represents firm-specific noise.³⁴ Given a linear contract of slope θ_{it} and normalizing $g = 1$, the CEO exerts effort $a_{it} = \theta_{it} b_i^*$, and the optimal slope is³⁵

$$\theta_{it} = \frac{1}{1 + \eta \left(\frac{\sigma_{FIRM}}{b_{it}} \right)^2}$$

Because investors set contracts based on perceived skill $\tilde{b}_{it}$ rather than b_i^* , systematic misperception produces systematic deviations in incentive provision. CEO i 's true skill, b_i^* , is assumed to remain constant over time.

Given the observed contract slope θ_{it} and candidate σ_{FIRM} , the model implies an invertible mapping³⁶

$$\tilde{b}_{it} = \sqrt{\frac{\eta \sigma_{FIRM}^2 \theta_{it}}{1 - \theta_{it}}} \quad (13)$$

This allows the empirical procedure to infer investors' perceived skill directly from compensation data.

5.3 Data

The structural estimation uses executive compensation and firm characteristics from WRDS ExecuComp from 2000–2025. The sample includes approximately 2,600 distinct firms comprising 3,693 firm-CEO spells, with an average of roughly eight observations per spell. Following the theory, the key objects constructed from the data are:

- $\theta_{it} = \frac{\text{variable compensation}}{\text{market capitalization}}$, the ratio of the CEO's performance-based pay granted in fiscal year t to the firm's market capitalization in the same year;³⁷
- $c_{it}(V_{it})$, estimated as the CEO's total compensation, ExecuComp's *tdc1* statistic;
- $\Delta V_{it} = \text{mktcap}_{it} - \text{mktcap}_{i,t-1}$, the annual change in firm value;³⁸ and
- $\tilde{b}_{it}$, the implied investor-perceived skill (from eq. (13)).

Firms with fewer than four observations are excluded from the estimation panel to ensure identification of CEO-by-firm-level b_i^* .

³⁴The baseline takes firm-value change due to fundamentals (ϵ_{it}^{FIRM}) to be mean-zero conditional on CEO skill. As a robustness check, I re-estimate the structural model with the firm-side shock centered at the sample-average growth rate. This allows firm value to be driven by secular drift rather than CEO skill and randomness alone. The results are qualitatively similar to the baseline specification's dispersion estimates and, if anything, strengthen the estimated average investor optimism ($\hat{\mu}_{INV}$ rises from 85 to 112; see Appendix E.3).

³⁵Derived from the standard CARA/Gaussian linear-contract environment with effort cost $\frac{\eta}{2} a^2$ and risk aversion η . I assume all CEOs share the same $\eta = 4 \times 10^{-4}$, which Yan et al. (2025) identify as a reasonable empirical calibration for CARA preferences.

³⁶In implementation, σ_{FIRM} enters the contracting formulas as the firm's dollar volatility $\psi_i \sigma_{FIRM}$, where ψ_i scales per-unit-of-size volatility by firm size. The scaling is introduced with the likelihood construction below.

³⁷This is a flow measure of incentive intensity in the Jensen–Murphy tradition. The dollar value of performance-linked pay granted in the year, scaled by firm value. A stock-based alternative, the delta of the CEO's accumulated equity and option portfolio, would capture wealth-performance sensitivity but would require option-grant detail and valuation assumptions outside the present data. The flow measure matches the model's timing (the shareholder base sets each year's contract) and is the margin say-on-pay votes ratify. Re-estimation under a portfolio-delta measure of θ is an intuitive robustness extension and is left for future work.

³⁸A potential concern with this specification is that the change in market capitalization may conflate value creation with net issuance and buybacks. Re-estimating the structural model with value creation measured instead as the firm's annual total shareholder return times lagged capitalization (built from CRSP total returns so that performance is measured from returns actually delivered to shareholders) yields parameter estimates within the baseline's bootstrap confidence intervals (See Appendix E.3).

5.4 Likelihood Estimation

The likelihood consists of two components: (i) the cross-sectional distribution of investor misperception and (ii) the distribution of firm performance.

Component 1: Investor Beliefs—For firm i , misperceptions are:

$$\begin{aligned}\epsilon_{it}^{INV} &= \tilde{b}_{it} - b_i^*, \\ \epsilon_{it}^{INV} &\sim \text{Normal}(\mu_{INV}, \sigma_{INV}^2)\end{aligned}$$

For each candidate vector of population parameters $(\sigma_{INV}, \sigma_{FIRM}, \mu_{INV})$, each firm's b_i^* is concentrated out in an inner optimization that maximizes the firm's combined likelihood across both the investor belief likelihood component and the firm performance likelihood component as described in the “Firm-Specific Skill” paragraph below. The outer optimizer therefore maximizes a profile likelihood over the three population parameters. Notably, the functional form imposed on investor misperception allows the structural estimation to recover aggregate bias in either direction. The mean μ_{INV} is a free parameter, so pessimistic shareholder bases are admissible, and the systematic overvaluation predicted by the adverse-selection mechanism becomes a testable restriction ($\mu_{INV} > 0$) rather than a maintained assumption. The corresponding log-likelihood contribution is:

$$\ell_i^{(1)} = \sum_t \log f_{\text{Normal}}(\epsilon_{it}^{INV}; \mu_{INV}, \sigma_{INV}^2)$$

The specification of the investor belief errors merits discussion because the theory of Section 3.2.6 provides a density equation for the pivotal shareholder's belief—an order statistic of the signal draws—though the estimation deliberately utilizes a related but slightly different formulation. That density is indexed by the environmental parameters of each firm's market, the number of potential investors N_{it} and the effective float s_{it} . Conditioning the likelihood estimates on these values would introduce practical difficulties—requiring thousands of non-verifiable calibrations, especially as to N_{it} —and, even worse, would assume the paper's conclusion by imposing a direction of investor bias rather than testing it. Furthermore, specifying $\epsilon_{it}^{INV} \sim \text{Normal}(\mu_{INV}, \sigma_{INV}^2)$ is without much loss of fidelity as central order statistics are asymptotically normal centered at the corresponding quantile of the signal distribution. The estimated $(\hat{\mu}_{INV}, \hat{\sigma}_{INV})$ are, in this sense, the identifiable reduced-form coordinates of the order-statistic model, and the Gaussian likelihood estimates them consistently. In addition, Appendix E.3 addresses the investor error distribution question directly with one-sided and log-normal alternatives; the sign and interpretation of the estimated distortion are invariant across specifications.

Component 2: Firm Performance—Given the model, we have:

$$\begin{aligned}\epsilon_{it}^{FIRM} &= \Delta V_{it} - \theta_{it}(b_i^*)^2 + c_{it}(V_{it}), \\ \epsilon_{it}^{FIRM} &\sim \text{Normal}(0, \sigma_{FIRM}^2)\end{aligned}$$

As the contract model abstracted away from variation in firm size, σ_{FIRM}^2 above does not scale well—it would imply the same dollar variance for a \$500 million firm as for a \$500 billion one. To address this issue, I adjust the fundamental shock by firm size. Specifically, in place of the second equation above I use $\epsilon_{it}^{FIRM} \sim \text{Normal}(0, \psi_i^2 \sigma_{FIRM}^2)$ where ψ_i is the median in-sample observation of firm i 's market capitalization and σ_{FIRM} is fundamental volatility per unit of firm size, yielding:

$$\ell_i^{(2)} = \sum_t \log f_{\text{Normal}}(\epsilon_{it}^{FIRM}; 0, \psi_i^2 \sigma_{FIRM}^2)$$

The same scaling applies wherever fundamental volatility enters the contracting formulas. The optimal slope and the inversion of Equation (13) are evaluated at the firm's dollar volatility $\psi_i \sigma_{FIRM}$, so that in implementation $\tilde{b}_{it} = \sqrt{\eta \psi_i^2 \sigma_{FIRM}^2 \theta_{it} / (1 - \theta_{it})}$. As robustness checks, in Appendix E I also allow ϵ_{it}^{FIRM} to have positive-mean distributions—reflecting secular growth trends—and re-measure firm performance by total shareholder return (which excludes net issuance) rather than the change in market capitalization. In these checks, the parameter estimates remain within their baseline confidence intervals and the adverse-selection finding is unchanged.

Firm-Specific Skill: For each candidate $(\sigma_{INV}, \sigma_{FIRM}, \mu_{INV})$, the firm-specific time-invariant CEO skill, b_i^* , is obtained by solving:

$$b_i^* = \arg \max_{b_i} \{\ell_i^{(1)} + \ell_i^{(2)}\}$$

Of note, the belief component alone cannot locate b_i^* because a level shift in candidate skill may be absorbed one-for-one by the estimated misperceptions. Hence, the performance component $\ell_i^{(2)}$ anchors the estimate by penalizing skill levels that rationalize contract history but are inconsistent with realized firm performance.

Global Parameters: Concentrating out b_i^* pins down a panel of $[\epsilon_{it}^{INV}, \epsilon_{it}^{FIRM}]$ for the candidate parameter vector, allowing computation of the full-sample likelihood as:

$$\mathcal{L}(\sigma_{INV}, \sigma_{FIRM}, \mu_{INV}) = \sum_{i=1}^{N_{\text{firms}}} \max_{b_i^*} \{\ell_i^{(1)} + \ell_i^{(2)}\}$$

I maximize $\mathcal{L}$ using a derivative-free Nelder–Mead optimizer. For the main specification the objective is smooth and the choice is one of convenience,³⁹ but the same optimizer is used for the constrained robustness specifications whose inner-loop boundary makes the objective non-smooth so that estimates are comparable across specifications.

³⁹A multistart check from widely separated initializations, reported in Appendix E, converges to a single optimum.

Post-Estimation Quantities: From the likelihood-maximized parameters $(\hat{\sigma}_{INV}, \hat{\sigma}_{FIRM}, \hat{\mu}_{INV})$ and corresponding firm-specific $\hat{b}_i^*$, I compute the estimated firm-year belief errors:

$$\hat{\epsilon}_{it}^{INV} = \tilde{b}_{it} - \hat{b}_i^*$$

These estimated misperceptions underlie the economic-magnitude accounting in Section 5.5 and the sorting variable of the asset-pricing exercise in Section 6.2. I also compute each firm-year’s expected value net of compensation under undistorted beliefs, $\mathbb{E}[V(a(b^*), \epsilon^{FIRM}) - c(V(a(b^*), \epsilon^{FIRM}))]$ evaluated at $\hat{b}_i^*$. Comparison of this with the same quantity at distorted beliefs yields the “cost of investor adverse selection” of Section 5.5.

5.5 Results

Table 2 reports the maximum-likelihood estimates with bootstrap confidence intervals. Investor noise as to CEO skill ($\hat{\sigma}_{INV} = 1,986$) is substantially smaller than scaled firm financial noise ($\hat{\sigma}_{FIRM} = 0.538$, or roughly $\hat{\sigma}_{FIRM} \psi_i \approx 1.2 \times 10^6$ at the median firm). Which is to say, most of the variation in firm outcomes is fundamental, and the belief errors the model recovers are a comparatively small disturbance layered on top of it.

Table 2: Structural parameter estimates (N μ specification)

Parameter	Estimate	95% CI
σ_I	1,986	[991, 3,220]
σ_F	0.538	[0.364, 0.769]
μ_I	85	[47, 137]

Notes: Maximum-likelihood estimates with 95% confidence intervals from 1,000 cluster bootstrap replications (firms resampled with replacement, full re-estimation per replication). Pairs-bootstrap intervals (firm-years resampled) are similar throughout, agreeing within 5–15%.

The main parameter of interest is the mean of investor errors. The estimate is positive and statistically distinguishable from zero: $\hat{\mu}_{INV} = 85.2$, with a 95% confidence interval of [47, 137] under the cluster bootstrap and [51, 152] under the pairs bootstrap. While the specification permitted misperception of either sign—a shareholder base that was on average correct ($\mu_{INV} = 0$) or pessimistic ($\mu_{INV} < 0$)—the data recovered an estimate of systematic overvaluation ($\mathbb{P}(\hat{\mu}_{INV} > 0) = 0.999$).

The magnitude of $\hat{\mu}_{INV}$ is also informative and economically meaningful—since the median $\hat{b}^*$ is 760, the average shareholder base perceives their CEO as $\approx 11\%$ more productive than they truly are. The extra annual value creation expected is, accordingly, 0.4 basis points of firm value per year at the median firm (or ≈ 4.5 bps at the mean). This optimism translates into steepening the median CEO’s contract by enough to incur roughly \$0.5 million per year of additional incentive compensation (about twelve percent of median total pay, in 2024 dollars). Of this, however, the bulk is offset by the additional effort the steeper contract elicits. It is the dispersion of belief errors (the larger magnitude $\hat{\sigma}_{INV}$) that generates the majority of the estimated \$1.16 million average annual loss (see Section 5.5) to shareholders from CEO-skill misperception.

Estimates of investor errors, ϵ_{it}^{INV} , which are reflected in the “costs” depicted in Figure 2, reveal significant dispersion in the amount of investor misperception by firm. Most firms exhibit little-to-no investor-induced CEO skill mispricing. The median annual loss per firm is around \$67,000 in 2024 dollars. The mean loss, at \$1.16 million, is more than an order of magnitude larger, indicating that the aggregate cost is driven by a tail of severely misperceived firms. This represents about 16% of aggregate CEO compensation in the sample. The top decile of firm-years accounts for approximately 85% of the aggregate loss. The top one percent alone accounts for about a third of the total. To put the central tendency in context, the median loss represents about $2.9 \times 10^{-3}\%$ of median market capitalization (the mean loss is roughly 0.009% of mean market capitalization), indicating that these costs are unlikely on their own to be a major driver of stock performance for most companies—though consistent identification of the outliers may hold potential to do so.

Figure 2 and Table 3 report the estimated annual dollar cost of misperception-driven CEO overcompensation from 2002–2023.⁴⁰ Summed over the sample, estimated losses total roughly \$34 billion in 2024 dollars across approximately 29,000 firm-years. The time pattern aligns with the mechanism. Per-firm costs peak during years of major valuation dislocation. The year 2008 is the costliest in the sample with a mean loss of \$3.1 million per firm. The bootstrap confidence intervals in Table 3 are wide, reflecting the outlier-driven skew documented above. However, the sign is unambiguous. The cluster-bootstrap confidence intervals for both the mean and the median loss lie entirely below zero in all 22 years of the sample.

The composition of the loss is as important as its size. I decompose the total by counterfactually re-contracting under beliefs stripped of their idiosyncratic noise (retaining only the systematic bias $\hat{\mu}_{INV}$) and, conversely, under beliefs stripped of the bias (retaining only the noise). The systematic adverse-selection component accounts for roughly \$0.7 billion of the \$33.6 billion total,⁴¹ or about two percent. Idiosyncratic belief dispersion accounts for nearly ninety percent. The small remainder is attributable to the interaction of the two.⁴² This finding informs the policy discussion in Section 7 in that governance mechanisms that correct average beliefs target the smaller component of the loss.

6 Empirical Context for the Structural Estimation Results

6.1 Do Shareholders Vote Against Misaligned Pay? Evidence from Say-on-Pay

It is often remarked that, broadly put, shareholders’ corporate governance options boil down to “sell, vote, sue.” So far, this paper has discussed the *Tornetta* litigation and market allocation dynamics, leaving “vote” as the unexplored margin. I turn to that now.

⁴⁰End dates were truncated due to incomplete data in 2000 and 2025 and the lagged nature of contract realization (θ_{it} is only computable in $t + 1$ once the dollar amount of the final payout is observed).

⁴¹Aggregate figures are totals over sample firms—roughly the ExecuComp universe after data linkages, averaging 1,288 firms per year (2002–2023). This is about a fifth of CRSP securities by count but 54% of listed market value on average (60% by the end of the sample); because both pay and the estimated losses are concentrated in large firms (Section 7), the in-sample aggregates capture the bulk of the dollar phenomenon and are conservative as economy-wide figures.

⁴²The components do not sum exactly to the total because the cost is convex in the belief error: noise centered at $\hat{\mu}_{INV}$ is more expensive than noise centered at zero. Bootstrap intervals: selection component [0.13, 1.07] billion; noise component [8.4, 40.4] billion.




Figure 2: Concentration of misperception losses across firm-years. The curve plots the cumulative share of the aggregate estimated loss (2024 dollars) against the share of firm-years, sorted from the largest loss. The dotted diagonal is the equal-losses benchmark. Losses are the same winsorized per-observation values used by the cost tables, so the displayed concentration is conservative: winsorization clips the extreme tail. The Gini coefficient of the loss distribution is 0.89.

The Dodd-Frank say-on-pay regime provides for an advisory shareholder vote on executive compensation contract, which is of obvious relevance to the topic at hand. I construct a dataset of say-on-pay votes from SEC filings. Vote tabulations are reported in Item 5.07 of Form 8-K. I collect all such filings for the estimation-sample firms (23,061 filings, 2010–2026) and extract the vote counts on the management say-on-pay proposal, yielding 19,456 firm-year votes.⁴³ I calculate an “approval” metric as the share of votes cast in favor: i.e., $\text{For}/(\text{For}+\text{Against})$. The median vote passes with 95.5% approval, but failures occur in every year of the sample. Each vote in calendar year t is matched to the structural estimates and financials of fiscal year $t - 1$, the year whose compensation is being ratified.

Figure 3 presents the main findings of this investigation. Approval is not monotone in estimated misperception. It peaks almost exactly at the model’s zero-misperception point and declines in both directions. Shareholders dissent from both contracts that overpay relative to inferred skill and from those that underpay. Table 4 quantifies these dynamics. A one standard deviation increase in the distance of misperception from zero reduces approval by 1.9 percentage points ($t = -5.5$; column (4)) and raises the probability of a below-70% vote by 3.2 percentage points ($t = +5.3$; column

⁴³Extraction combines a rule-based parser with a large language model applied to the archived filing text, each serving as an independent check on the other; exact agreement on both vote counts is 76.1% where both methods extract, and all disagreements, near-failing votes, and formatting edge cases were adjudicated against the underlying filings. Appendix D.1 describes the pipeline. The resulting failure incidence — 1.9% of votes below majority support and 6.3% below the 70% threshold that triggers proxy-advisor responsiveness review — matches rates reported for the say-on-pay universe, an external check on the extraction.

Table 3: Yearly cost of investor adverse selection (2024\$M)

Year	N	Mean	95% CI	Median	95% CI	Total
2002	827	-1.85	[-4.35, -0.76]	-0.088	[-0.197, -0.034]	-1,530
2003	972	-1.05	[-2.45, -0.42]	-0.034	[-0.075, -0.015]	-1,020
2004	1,102	-0.96	[-2.27, -0.38]	-0.039	[-0.084, -0.016]	-1,060
2005	1,143	-0.98	[-2.22, -0.40]	-0.049	[-0.106, -0.020]	-1,125
2006	1,089	-0.96	[-2.17, -0.40]	-0.054	[-0.117, -0.023]	-1,049
2007	1,061	-1.16	[-2.57, -0.46]	-0.057	[-0.131, -0.025]	-1,228
2008	1,128	-3.13	[-6.89, -1.33]	-0.191	[-0.434, -0.077]	-3,533
2009	1,479	-0.93	[-2.10, -0.38]	-0.049	[-0.108, -0.021]	-1,382
2010	1,551	-0.87	[-1.91, -0.35]	-0.044	[-0.098, -0.018]	-1,352
2011	1,570	-1.33	[-3.07, -0.54]	-0.069	[-0.155, -0.029]	-2,087
2012	1,543	-0.89	[-2.01, -0.37]	-0.060	[-0.132, -0.024]	-1,377
2013	1,521	-0.62	[-1.37, -0.26]	-0.044	[-0.098, -0.019]	-950
2014	1,455	-0.80	[-1.81, -0.32]	-0.047	[-0.102, -0.020]	-1,166
2015	1,424	-0.93	[-2.10, -0.39]	-0.068	[-0.149, -0.029]	-1,322
2016	1,407	-0.70	[-1.52, -0.29]	-0.058	[-0.125, -0.024]	-982
2017	1,421	-0.88	[-1.95, -0.36]	-0.060	[-0.127, -0.026]	-1,257
2018	1,428	-1.51	[-3.33, -0.62]	-0.101	[-0.227, -0.043]	-2,155
2019	1,413	-1.12	[-2.51, -0.46]	-0.063	[-0.134, -0.027]	-1,584
2020	1,421	-1.32	[-2.97, -0.54]	-0.084	[-0.184, -0.036]	-1,878
2021	1,433	-1.14	[-2.53, -0.47]	-0.085	[-0.183, -0.036]	-1,629
2022	1,355	-1.64	[-3.67, -0.69]	-0.173	[-0.373, -0.074]	-2,219
2023	1,268	-1.37	[-3.03, -0.58]	-0.162	[-0.340, -0.069]	-1,732
Overall	29,011	-1.16	[-2.51, -0.50]	-0.067	[-0.142, -0.030]	-33,617

Notes: Annual cost of misperception-driven CEO overcompensation, in 2024 dollars (millions), under the $N\mu$ specification with the corrected contract-intercept formula: point estimates with 95% cluster-bootstrap confidence intervals (1,000 replications). Costs are winsorized at 1% pooled across years. Total is the aggregate across all sample firms in the year.

(6)) against a 6% base rate, conditional on excess pay, size, performance, and year and industry effects. The signed response survives inclusion of firm fixed effects (column (2)). This two-sided pattern is a distinctive implication of the model. As preferences over contracts are single-peaked in perceived skill, dissent should rise with misalignment in either direction.

One initial question about Figure 3 is whether shareholder votes simply respond to observed contract steepness. That is, if θ is too high or too low, a subset of shareholders might vote “no” reflexively, meaning Figure 3 merely recovers the fact that some contracts are steep and others are flat. To address this, I orthogonalize the $\hat{\epsilon}_{INV}$ s to contract steepness and regress Say-on-Pay approval rates on the absolute value of those residuals (i.e., on $|\tilde{\epsilon}^\perp|$, see Table 4 column (5)). The result holds in magnitude and significance (−1.9 percentage points per SD, $t = -5.7$; Table 4, column (5)); see also Appendix D.1.1).

Another foreseeable question is why votes should track a component of misperception that Section 3.2.6 says no participant can precisely identify. The answer is that shareholder objections to compensation packages do not require precise identification. If dissent is triggered by distance of a pay package from a voter’s ideal point, then what matters is the distance of the median shareholder’s estimation from the marginal shareholder, not any given shareholder’s ability to invert the CEO’s true skill. Larger pivotal errors—in either direction—plausibly accompany greater dispersion between the pivotal belief and those of the rest of the base,⁴⁴ so approval falls with $|\hat{\epsilon}^{INV}|$ even though no voter sees

⁴⁴Formally, this is a conjecture about conditional order-statistic spacings, which I leave for future work.




Figure 3: **Say-on-pay approval and estimated misperception.** Each point is one of 25 quantile bins of firm-year votes ($n = 13,077$), plotting the mean approval share For/(For+Against) against mean estimated misperception $\hat{\epsilon}_{i,t-1}^{INV}/(\hat{\sigma}_{FIRM}\psi_i)$ in standard-deviation units; vertical bars are 95% confidence intervals. Approval is residualized on excess pay, log market capitalization, prior-year firm value growth, and year and industry fixed effects (grand mean added back). The dashed line is the piecewise-linear fit with a kink at zero misperception; the dotted vertical line marks zero. Standard errors clustered by firm.

Table 4: Say-on-pay approval and estimated misperception

	(1)	(2)	(3)	(4)	(5)	(6)
Misperception $\tilde{\epsilon}$ (SD)	-0.0085*** (-3.74)	-0.0113*** (-4.62)				
over-perception side			-0.0177*** (-5.08)			
under-perception side			0.0311*** (3.85)			
Distance $ \tilde{\epsilon} $ (SD)				-0.0192*** (-5.52)		0.0319*** (5.29)
Distance, benchmark comp. $ \tilde{\epsilon}^\perp $ (SD)					-0.0185*** (-5.65)	
Excess pay	-0.0285*** (-7.90)	-0.0218*** (-5.55)	-0.0287*** (-8.15)	-0.0272*** (-8.64)	-0.0307*** (-9.66)	0.0373*** (6.72)
Controls	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Industry FE	✓		✓	✓	✓	✓
Firm FE		✓				
Observations	13,077	13,077	13,077	13,077	13,077	13,077
R^2	0.066	0.350	0.077	0.076	0.079	0.042

Notes: Dependent variable: say-on-pay approval share For/(For+Against) in columns (1)–(5); indicator for approval below 70% in column (6). Misperception is the N μ -specification $\hat{\epsilon}_{i,t-1}^{INV}/(\hat{\sigma}_{FIRM}\psi_i)$, standardized; the vote in year t concerns fiscal year $t - 1$ compensation. Column (3): separate slopes on the positive and negative parts of misperception (kink at zero). Column (5): distance in the component of misperception orthogonalized to the contract slope $\hat{\theta}$ within year (the performance-benchmark component; Table 9 reports the full decomposition), additionally controlling for absolute firm value growth. Controls: excess pay (residual of log total compensation on log market capitalization, year, and industry), log market capitalization, prior-year firm value growth (winsorized 1%). t -statistics in parentheses, standard errors clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the error itself. The correlation reflects shared exposure to the same draw of beliefs, not inference from it.




Figure 4: **The vote data independently locate the model’s alignment point.** Residual sum of squares from re-estimating the piecewise say-on-pay approval regression with the kink placed at candidate locations c (in standard deviations of estimated misperception). The fit-maximizing kink is $\hat{c} = +0.10$ SD, with firm-clustered bootstrap 95% confidence interval $[-0.31, +0.24]$ (shaded band). The band contains the model’s zero-misperception point ($c = 0$, dotted) and excludes the sample mean of the sorting variable ($+0.34$ SD, dashed), where a firm-typical-slope alternative would place the approval peak. Approval residualized on excess pay, size, prior-year growth, and year and industry fixed effects.

The voting evidence doubles as an out-of-sample validation of the structural estimates. The vote data share no inputs with the estimation—votes are tabulated from 8-K filings, while $\hat{\varepsilon}^{INV}$ is constructed from compensation contracts and firm performance—yet shareholder approval peaks almost exactly at the model’s zero-misperception point. Re-estimating the piecewise specification with the kink at candidate locations, the data select $\hat{c} = +0.10$ standard deviations (firm-clustered bootstrap 95% confidence interval $[-0.31, +0.24]$). This interval contains the model’s zero and excludes the sample mean of the sorting variable, where a “firms are punished for deviating from their typical contract” alternative would place the peak (Figure 4). Likewise, the model’s misalignment distance drives out naive benchmarks based on the firm’s own slope history or industry norms when raced jointly ($t = -4.6$ against $t = -1.7$ and a sign flip, respectively), despite correlations near 0.75 (see Appendix D.1.2 for an expanded discussion).

6.2 Empirical Asset Pricing Test

Section 5.5 reported that the total effect of CEO skill misperception is small compared to the market capitalization of most firms. This modesty relative to firm value is expected, as the direct “harm,” within the context of the model, to investors is capped by the amount of CEO pay as extra incentivization should simply lead to more (if inefficient) value-creating effort from the CEO. That said, it is possible that the direct cost may not be the whole story. If—as per Bebchuk et al. (2011b)—excessive CEO pay is due to managerial rent extraction, then overly optimistic shareholder

estimation of the CEO may imply lax oversight of management in other areas, becoming, in a sense, a canary in the governance coal mine.⁴⁵ In other words, the question is whether this paper’s “misperception” correlates with wider failures of firm governance.

This subsection runs an asset pricing test to evaluate that question. The test sorts on the standardized firm-year belief error, $\hat{\epsilon}_{it}^{INV}/(\hat{\sigma}_{FIRM} \psi_i)$, which scales estimated misperception by the firm’s fundamental volatility so the sort is not mechanically a sort on firm size. At the end of each calendar year t , all firms with a non-missing estimate are sorted into quintiles. Portfolio membership is assigned in July of year $t + 1$ and held for twelve months without rebalancing. Monthly returns are from CRSP, value-weighted by lagged market capitalization. The sample spans July 2002 through December 2023, for a total of 258 monthly observations. Table 5 reports Fama–French five-factor alphas for the five quintile portfolios, with 95% confidence intervals from 1,000 cluster-bootstrap replications in which the structural model is re-estimated and the portfolios re-formed from scratch.

Table 5: Quintile portfolio alphas: optimistically held firms earn ordinary returns

Portfolio	α	(t)	Bootstrap 95% CI	EW α	(t)
Q1 (most under-perceived)	-0.0017	(-1.53)	[-0.0044, 0.0003]	-0.0002	(-0.25)
Q2	-0.0010*	(-1.84)	[-0.0023, 0.0004]	-0.0004	(-0.56)
Q3 (aligned)	0.0022***	(2.88)	[0.0007, 0.0040]	0.0014	(1.38)
Q4	0.0042***	(3.80)	[0.0015, 0.0060]	0.0006	(0.34)
Q5 (most over-perceived)	0.0023	(1.10)	[0.0003, 0.0052]	0.0023	(1.34)
V: long Q3, short $\frac{1}{2}(Q1+Q5)$	0.0019	(1.27)	[0.0000, 0.0040]	0.0003	(0.20)
Directional: Q1–Q5	-0.0040*	(-1.79)	[-0.0088, -0.0008]	-0.0026	(-1.34)

Notes: Monthly FF5 alphas of value-weighted portfolios formed on $\hat{\epsilon}_{t-1}^{INV}/(\hat{\sigma}_{FIRM} \psi_i)$ quintiles (annual sort, July–June holding). HAC(12) t -statistics in parentheses. Square brackets: 95% confidence intervals from 1,000 cluster-bootstrap replications in which the structural model is re-estimated and portfolios re-formed. Stars mark leg-level HAC t -statistics and are descriptive: with five quintiles per sort, individual legs invite multiple-comparison false positives, so joint inference is carried by the Gibbons–Ross–Shanken test, which rejects all five alphas being zero in the full sample ($F_{5,247} = 5.54$, $p = 0.0001$) but not in the expanding-window construction ($F_{5,187} = 0.72$, $p = 0.61$). The V portfolio tests whether misalignment *per se* carries a return premium (it does not; see also Table 6); the directional spread is reported as an open pattern and examined in Appendix E.1.1. Final columns: expanding-window (real-time) construction.

The full-sample results run against the canary story. The starred alphas sit in the aligned and mildly over-perceived quintiles: in the latter case, significantly positive, where the “proxy for more extraction” story would predict a negative coefficient. The point estimate of most over-perceived quintile even indicates earnings of +0.23% per month. In the expanding-window construction, which re-estimates the model each year using only data available at portfolio formation, every significance star drops out, and the joint test cannot reject all five alphas being zero ($F_{5,187} = 0.72$, $p = 0.61$, versus $F_{5,247} = 5.54$, $p = 0.0001$ in the full sample). Furthermore, translating the canary story into a portfolio strategy—that an investor base who most over-perceives their CEO will have the “loosest reins” on managerial rent extraction, but those who most under-perceive their CEO will have the tightest ones (so Q1–Q5)—if anything loses money (−0.40% ($t = -1.79$) or, with the expanding window estimation, −0.26% ($t = -1.34$)).

A second hypothesis might be that any misalignment destroys value; “too tight” reins are just as bad as “too loose” ones. The intuitive test for this is the symmetric “V” portfolio that is long the aligned middle quintile and short one-half

⁴⁵Other anecdotes suggest that excessive deference to charismatic CEOs may result in detrimental firm outcomes. See Fan and Nguyen (2025) arguing that a charismatic founder, Sam Bankman-Fried, used investor esteem to relax standard oversight in order to relax governance at FTX.

each of the two extremes. Table 6 reports the results of that estimation. The simple-sort V earns +0.19% per month ($t = 1.27$; bootstrap interval [0.00%, +0.40%]), the size-balanced double-sort version earns -0.19% ($t = -1.95$), but both of these effects vanish in the expanding-window constructions ($t = 0.20$ and $t = -0.99$). Point estimates that flip sign across constructions and vanish in real time are indicia of a lack of stable premium. The factor loadings, reported in Table 6, show why the raw and risk-adjusted comparisons differ. The V portfolio is systematically anti-market and anti-size. Firms at the misperception extremes are higher-beta and smaller than the aligned middle, which might make sense from an investor-noise perspective if smaller firms are more information-scarce, leading to higher noise. These exposures are stable across constructions, including in real time, and once they are absorbed no construction leaves an abnormal return.

Table 6: The V portfolio: misalignment *per se* is not priced

	Full sample		Expanding window	
	Simple	Double sort	Simple	Double sort
α (monthly)	0.0019 (1.27)	-0.0019* (-1.95)	0.0003 (0.20)	-0.0015 (-0.99)
MKT	-0.127*** (-3.75)	-0.147*** (-6.40)	-0.148*** (-3.83)	-0.140*** (-4.90)
SMB	-0.221*** (-3.20)	-0.154*** (-3.34)	-0.053 (-1.01)	-0.084 (-1.52)
HML	0.154* (1.94)	0.038 (0.54)	0.052 (0.94)	0.024 (0.38)
UMD	0.063 (1.03)	0.040 (1.12)	0.004 (0.06)	0.113*** (2.71)
RMW	-0.028 (-0.41)	0.075 (1.47)	0.115 (1.45)	0.038 (0.52)
CMA	0.158 (1.32)	0.121 (1.47)	0.266*** (2.66)	0.269*** (2.88)
Months	258	258	198	198
R^2	0.267	0.340	0.268	0.403

Notes: Fama–French five-factor regressions of the symmetric “V” portfolio: long the aligned middle misperception quintile, short one-half each of the extreme quintiles, value-weighted, formed annually (July–June holding). The double sort forms misperception sub-quintiles within size quintiles and pools across size buckets, so the V is size-balanced by construction. Self-financing portfolio: alphas estimated on raw spread returns. HAC(12) t -statistics in parentheses; stars are descriptive. Expanding-window columns re-estimate the full structural model using only data through each formation year. The full-sample simple-sort V carries a cluster-bootstrap 95% confidence interval of [0.00%, 0.40%] (1,000 replications, structural model re-estimated per replication).

For completeness, the quintile alphas do vary jointly (the Gibbons–Ross–Shanken test rejects all five being zero, $F_{5,247} = 5.54$, $p < 0.001$), and the directional Q1–Q5 spread earns -0.40% per month ($t = -1.79$) in the full-sample construction while remaining statistically indistinguishable from zero in real-time constructions. I report this directional pattern as an open question rather than build on it. Appendix E.1.1 presents the full battery and some candidate interpretations.

While investor misperception of CEO ability, whose direct impact is capped by the total amount of CEO pay, shows no robust penalty to risk-adjusted returns, this optimism leaves one governance footprint outside of pay. CEOs of the most over-perceived firms are less likely to depart than those of aligned firms (4.2 percentage points per year against a 7.5 percent base rate (firm-clustered $t = -2.37$); Appendix D.5).

7 Counterfactuals and Discussion

The results so far describe a corporate governance failure with an unusual structure. The shareholder base acting rationally, albeit with a noisy information set, creates a real and measurable distortion to firm value. If public markets are inherently informationally noisy,⁴⁶ then the question becomes whether policy intervention is inherently futile or if opportunities for improvement exist.

If a workable intervention exists, some actor would need to implement it. So, let's first consider whether there exists any agent who is better situated to directly set CEO compensation contracts. Any such candidate would have to have a better assessment of CEO skill than the market forms on its own. While plausibly less-biased alternatives exist, each comes with potentially serious trade-offs. Within the context of this paper's model, CEOs know their true skill, though allowing them to set their own pay creates obvious agency problems. Independent directors may be recruited for subject-matter knowledge rather than enthusiasm for the company. They may, therefore, hold less-correlated signals as to CEO skill and, indeed, are even required by the major listing exchanges to comprise compensation committees under most circumstances.⁴⁷ However, the empirical results of Section 5.5 were obtained despite these listing standards, which may have to do with the director-responsiveness critique of Section 3.2.8. Thus any reforms leveraging director independence would have to further expand this already central role. A similar lack of selection pressure may also recommend further examination of the roles played by proxy advisors and passive institutional investors, who buy a slice of every firm. Nonetheless, what these actors gain in reduced directional bias, they plausibly lose in information aggregation capacity compared with the market's price mechanism.

One alternative to agent-based interventions is for a market-wide policy regime—i.e., “blunt instrument”—conditioned on observables. Perhaps the most direct (and bluntest) approach would be a cap on incentive pay, which might take the form of a uniform ceiling $\bar{\theta}$ on the contract slope. Of course, base pay would continue to be allowed to adjust through the contracting model at unchanged beliefs.⁴⁸ This intervention is, moreover, an implementable counterfactual given this paper's model and data. To run it, for each firm-year I replace the observed slope with $\min(\theta_{it}, \bar{\theta})$, evaluate the change in expected firm value at the estimated true skill $\hat{b}_i^*$ and aggregate. Figure 5 traces the resulting policy frontier across cap levels, with a 95% band from re-running the entire exercise inside each of the 1,000 bootstrap replications of the structural estimation.

This counterfactual study informs the policy question in a few ways. First, a uniform cap may create value. The recovered optimum sits at roughly the 58th percentile of observed slopes (bootstrap interval: 48th–62nd) and yields a net gain of about \$0.46 billion per year (95% interval [0.12, 0.64]). At this point, expected CEO pay falls by roughly \$1.8 billion per year—about a fifth of the aggregate pay bill. On the other hand, the flatter contracts induce lower CEO

⁴⁶See, e.g., Grossman and Stiglitz (1980); Allen (1990).

⁴⁷See O'Brien (2026).

⁴⁸The experiment is defined on the incentive slope rather than on dollar pay levels because the level of base pay is a free normalization in the model (ω is not identified). A cap on variable-pay intensity is the version of the policy the estimates can price. As in the theoretical discussion, if the CEO's participation constraint does not bind, a cap on variable pay effects a cap on total compensation.




Figure 5: Estimated welfare effects of a uniform cap on incentive pay. For each cap level $\bar{\theta}$ (expressed as a percentile of the observed distribution of contract slopes), each firm-year’s slope is replaced by $\min(\theta_{it}, \bar{\theta})$, base pay adjusts per the contracting model at unchanged beliefs, and the change in expected firm value is evaluated at the estimated true skill $\hat{b}_i^*$, aggregated in 2024 dollars per sample year. Dashed lines decompose the net effect into gains at over-incentivized firms and losses at firms whose estimated skill justified their slope. Shading: 95% band from 1,000 cluster-bootstrap replications of the full estimation. Beliefs and shareholder-base composition are held fixed under the policy (partial equilibrium); dollar pay levels are not identified by the model, so the experiment is defined on the incentive slope.

effort, resulting in \$1.3 billion in foregone value creation, netting the \$0.5 billion frontier gain. Second, the frontier is remarkably flat near its peak (see Figure 5). Any cap between roughly the 40th and 75th percentile captures most of the attainable gain, suggesting that a conservative cap at the higher end of this range is likely to capture most of the potential benefit and, in case imprecisely estimated, avoid the risk of net value destruction that the counterfactual identifies as occurring below roughly the 20th percentile. Third, even the optimal cap’s net recovery accounts for only about one third of the total distortion identified in Section 5.5 (\$0.46 billion of the \$1.53 billion per year that firm-by-firm correction of beliefs would deliver). Two reasons explain this. Since misperception can occur in either direction, the cap only addresses one arm of the problem. However, since misperception is skewed towards optimism, it accounts for the bulk of the problem. Adding an optimal “floor” (Appendix D.4) therefore does not add much beyond the cap alone (\$0.47 billion with the combined cap-and-floor versus \$0.46 billion per year for cap-only). More importantly, however, is that this cap addresses extreme values, not noise. Noisy contracts within “reasonable” values of incentive-pay steepness account for the other two thirds of value-destruction.

Two caveats merit discussion. First, this counterfactual analysis is an exercise in partial equilibrium. Bereskin and Cicero (2013) document “contagion” in CEO pay⁴⁹ and Edmans et al. (2017) argue that pay transparency itself can

⁴⁹Bereskin and Cicero (2013) examine changes in CEO pay following a 1995 change in Delaware law that allowed Delaware firms to adopt more robust defenses to hostile takeovers, finding a 33% pay increase for Delaware executives that non-Delaware rivals then matched.

strengthen peer-driven ratcheting effects. Taken together, these suggest that there would be a strong incentive to defect from incentive pay caps that are not universally enforced—which may be a tall order in the international market for managerial skill.⁵⁰ Second, the political economy of such a reform is unpromising. Within the model, a faithfully implemented cap leaves CEOs indifferent—the savings are certainty-equivalent-neutral—but shareholders, for their part, may be rationally apathetic given the small size of the aggregate distortion in comparison to the median firm’s market capitalization.⁵¹ Even absent organized opposition, which might materialize if executives fear overreach, no constituency holds a concentrated stake in pushing the reform.⁵²

The pay cap’s addressability ceiling raises the question of how to approach the remaining distortion. A decomposition of the total cost suggests that, instead of indirectly targeting realized selection effects by curtailing outlier incentive contracts (per the pay cap/floor policy above), an altruistic regulator might more profitably focus on reducing the upstream error that leads to these misaligned contracts in the first place: i.e., on improving the informational environment to reduce investor noise. Table 7 reports estimates of the value that various improvements along this dimension would have. Scaling down the estimated belief-error distribution and re-solving every contract shows that a fifteen percent reduction in dispersion creates as much value as the optimal cap. Halving the dispersion recovers \$1.18 billion per year—roughly two and a half times larger than the optimal, downstream incentive pay cap-and-floor.

Table 7: Welfare gains from reducing investor belief dispersion

Reduction in belief dispersion	Value recovered (\$B/yr)	Share of total cost
10%	0.32	21%
16%	0.48	32%
25%	0.71	46%
50%	1.18	77%
75%	1.45	95%
100%	1.53	100%

Notes: Each row scales every firm-year belief error by λ (dispersion and selection tilt together, consistent with winner’s-curse selection scaling with signal noise), re-solves each firm’s contract at the scaled beliefs, and prices the change with the same firm-year accounting as the cap frontier. Reduction = $1 - \lambda$. For comparison, the optimal uniform cap on incentive intensity recovers \$0.46 billion per year; a reduction in belief dispersion of 15% matches it.

Which policies can reduce belief dispersion and at what cost is a question that I leave to future work; though Table 7 indicates that it may be the more promising margin for regulatory intervention. A natural starting point might be the disclosure-based, information-forcing architecture of securities regulation. Mandatory disclosure might reasonably have the effect of compressing the distribution of investor beliefs around firm fundamentals.

These findings also bear on the governance question that opened this paper: how much evidentiary weight should

⁵⁰However, one implication of the decomposition above is worth emphasizing: reduction in incentive pay alone—even unilaterally by a single firm—would not necessarily put the company at a disadvantage in a competitive market. This is because the pay savings from an incentive cap (which then pushes more pay into the CEO’s base salary to satisfy their participation constraint) are certainty-equivalent-neutral. So, the effect of such a policy (and the source of its welfare gains) would be to remove only the risk premium and excess-effort compensation that a steep contract entails. Contagion would threaten these savings only through outside offers that, say, embed rents which the firm must match or reflect indexing to non-certainty-equivalent-adjusted total pay.

⁵¹About $2.9 \times 10^{-3}\%$.

⁵²Of course, there is precedent for national crises spurring policy creation in this area: both the \$500,000 cash-compensation cap for bailout recipients (See “In Curbing Pay, Obama Seeks to Alter Corporate Culture” *New York Times*, Feb. 4, 2009, by Stephen Labaton and Vikas Bajaj, www.nytimes.com, accessed December 1, 2025.) and the Dodd-Frank say-on-pay apparatus (Dodd-Frank Wall Street Reform and Consumer Protection Act, Pub. L. 111-203 (July 21, 2010).) emerged in response to the 2008 financial crisis.

the law place on shareholder approval? Ratification mechanisms premised on shareholder approval—from advisory votes to the cleansing framework of amended DGCL §144⁵³—inherit the informational boundary characterized in Section 3.2.6 and the adverse-selection phenomenon documented in Section 5.5. To the extent a cleansing vote concerns a question that is itself the object of investor adverse selection—perceptions of executive skill in the Tesla saga being the motivating example—approval by the minority may deserve less deference than current doctrine affords, or at least less than parity with the input of less adversely selected parties, such as independent directors.⁵⁴ A more ambitious application—weighting governance inputs by the informational exposure of the electorate—is worth stating even if this paper cannot give it a full assessment. Distinguishing among types of investors may be a slippery slope best left untouched. However, the theory implies that decision-makers could rationally give more weight to passive investors’ assessments of matters that plausibly motivated shareholders’ selection into the stock, and less weight where the issue does not implicate that selected assessment.⁵⁵

Taken together, the counterfactuals leave a clear agenda for future work. The full correction of investor beliefs is unattainable by the universe’s lack of perfect information. A uniform cap-and-floor policy captures only a minority of the distortion visible in CEO pay contract terms. Identifying ways in which to improve the informational environment in which shareholder bases form, therefore, holds the most promise for reducing the majority of the distortion this paper identifies. Short of that, however, reassessment of the role shareholder approval occupies in the hierarchies of governance mechanism is warranted.

8 Conclusion

This paper develops and tests a theory of how adverse selection among shareholders can distort executive compensation and its surrounding governance mechanisms. By embedding the Holmström–Milgrom contracting model into a market with heterogeneous investor signals, I show that, in expectation, optimistic investors disproportionately select into the shareholder base inducing a systematic overvaluation of CEO ability. Even when all agents behave rationally given their information, this selection effect pushes shareholders to approve incentive-heavy contracts that induce inefficiently high levels of effort. While such contracts boost observable performance, the gains do not offset the excess compensation costs, resulting in lower firm profits.

The paper’s empirical analysis quantifies this mechanism. Using compensation, ownership and performance data from 2000–2025, the structural estimation finds a positive and statistically significant mean investor belief error, corroborating the model’s adverse-selection prediction. In economic terms, these losses are meaningful in the context of CEO pay, but modest in terms of overall firm returns. I find this distortion results in mean losses of \$1.16 million per

⁵³See “What Newly Amended DGCL §144 Says (and Does Not Say) about Controlling Stockholder Transactions,” Marcel Kahan, available at corpgov.law.harvard.edu; posted June 5, 2025, accessed September 11, 2025.

⁵⁴For instance, this might inform how to decide a case in which an independent committee recommends against a pay package for a controlling CEO but a majority of the minority approves it.

⁵⁵An example of the latter: choosing the highest bidder in a sale of the firm—incentives are aligned on price regardless of investors’ initial assessments when buying in, though this may require additional assumptions as to voters’ strategic behavior.

firm per year (median \$67,000, in 2024 dollars). This represents only about $2.9 \times 10^{-3}\%$ of the median company's market capitalization, but roughly sixteen percent of aggregate CEO pay in the sample. An asset-pricing test rejects the natural "canary in the coal mine" reading of these distortions. In my sample, optimistically held firms do not appear to suffer any risk-adjusted return penalty, suggesting the measured contracting cost is not a proxy for larger, value-destroying governance issues.

This paper's counterfactual experiments explore the potential for policy intervention. A uniform cap on incentive intensity recovers about one third of the misperception-induced distortion. The counterfactuals show that the more promising margin for intervention is upstream in the process. As the contracting loss is convex in the belief error, reducing the dispersion of investor beliefs targets the problem in a more efficient way than policy interventions conditioned on realized belief errors can. A fifteen percent reduction in belief dispersion creates as much value as the optimal cap-and-floor policy; halving it recovers roughly two and a half times as much. Which policies can deliver such reductions in investor noise, and at what cost, is a question this paper poses for future work.

Together, these theoretical and empirical results qualify the role of shareholder approval as a governance safeguard. In informationally noisy environments, the shareholder base itself may be systematically biased in favor of excessively steep incentive contracts. In such environments, shareholder approval certifies that a contract's observable terms are unremarkable to the self-selected shareholder base, not that they are efficient.

References

- Allen, B. (1990). Information as an Economic Commodity [Publisher: American Economic Association]. *The American Economic Review*, 80(2), 268–273. Retrieved April 26, 2024, from <https://www.jstor.org/stable/2006582>
- Bebchuk, L. A., Cremers, K. J. M., & Peyer, U. C. (2011a). The CEO pay slice. *Journal of Financial Economics*, 102(1), 199–221.
- Bebchuk, L. A., Cremers, K. M., & Peyer, U. C. (2011b). The CEO pay slice. *Journal of Financial Economics*, 102(1), 199–221. <https://doi.org/10.1016/j.jfineco.2011.05.006>
- Bebchuk, L. A., & Fried, J. M. (2006). *Pay without performance: The unfulfilled promise of executive compensation* (First Harvard University Press paperback edition). Harvard University Press.
- Bereskin, F. L., & Cicero, D. C. (2013). CEO compensation contagion: Evidence from an exogenous shock. *Journal of Financial Economics*, 107(2), 477–493. <https://doi.org/10.1016/j.jfineco.2012.09.005>
- Bertrand, M., & Mullainathan, S. (2001). Are CEOs Rewarded for Luck? The Ones Without Principals Are*. *The Quarterly Journal of Economics*, 116(3), 901–932. <https://doi.org/10.1162/00335530152466269>
- Black, D. (1948). On the Rationale of Group Decision-making [Publisher: University of Chicago Press]. *Journal of Political Economy*, 56(1), 23–34. Retrieved February 18, 2026, from <https://www.jstor.org/stable/1825026>
- Cheng, I.-H., Hong, H., & Scheinkman, J. A. (2015). Yesterday’s Heroes: Compensation and Risk at Financial Firms [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12225>]. *The Journal of Finance*, 70(2), 839–879. <https://doi.org/10.1111/jofi.12225>
- Coles, J. L., Daniel, N. D., & Naveen, L. (2006). Managerial incentives and risk-taking. *Journal of Financial Economics*, 79(2), 431–468.
- Coles, J. L., Daniel, N. D., & Naveen, L. (2014). Co-opted boards. *The Review of Financial Studies*, 27(6), 1751–1796.
- Core, J. E., Holthausen, R. W., & Larcker, D. F. (1999). Corporate governance, chief executive officer compensation, and firm performance. *Journal of Financial Economics*, 51(3), 371–406. [https://doi.org/10.1016/S0304-405X\(98\)00058-0](https://doi.org/10.1016/S0304-405X(98)00058-0)
- David, H. A., & Nagaraja, H. N. (2003, July). *Order Statistics* (1st ed.). Wiley. <https://doi.org/10.1002/0471722162>
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets [Publisher: University of Chicago Press]. *Journal of Political Economy*, 98(4), 703–738. Retrieved November 24, 2025, from <https://www.jstor.org/stable/2937765>
- Edmans, A., & Gabaix, X. (2016). Executive Compensation: A Modern Primer. *Journal of Economic Literature*, 54(4), 1232–1287. <https://doi.org/10.1257/jel.20161153>
- Edmans, A., Gabaix, X., & Jenter, D. Executive Compensation: A Survey of Theory and Evidence We thank Aubrey Clark, Fabrizio Ferri, Carola Frydman, Tom Gosling, Steve Kaplan, Gaizka Ormazabal, Kelly Shue, Alexander Wagner, David Zhang, and especially Pierre Chaigneau for helpful comments and Janet Chater, Irem

- Erten, Jesus Gorrin, Joseph Kalmenovitz, and Jiaying Wei for research assistance. en. In: In *The Handbook of the Economics of Corporate Governance*. Vol. 1. Elsevier, 2017, pp. 383–539. ISBN: 978-0-444-63530-3. <https://doi.org/10.1016/bs.hecg.2017.11.010>
- Fahlenbrach, R. (2009). Shareholder Rights, Boards, and CEO Compensation. [Publisher: Oxford University Press / USA]. *Review of Finance*, 13(1), 81–113. <https://doi.org/10.1093/rof/rfn011>
- Fan, J. S., & Nguyen, X.-T. (2025). Founder Worship, Effective Altruism, and Corporate Governance. *Kentucky Law Journal*, 113(1), 1–70.
- Ferri, F., & Maber, D. A. (2013). Say on Pay Votes and CEO Compensation: Evidence from the UK. *Review of Finance*, 17(2), 527–563. <https://doi.org/10.1093/rof/rfs003>
- Fisch, J., Palia, D., & Solomon, S. D. (2018). IS SAY ON PAY ALL ABOUT PAY? THE IMPACT OF FIRM PERFORMANCE. *Harvard Business Law Review*, 8, 101.
- Gabaix, X., & Landier, A. (2008). Why Has CEO Pay Increased so Much? [Publisher: Oxford University Press]. *The Quarterly Journal of Economics*, 123(1), 49–100. Retrieved August 19, 2025, from <https://www.jstor.org/stable/25098894>
- Gayle, G.-L., Golan, L., & Miller, R. A. (2015). Promotion, Turnover, and Compensation in the Executive Labor Market [Publisher: [Wiley, The Econometric Society]]. *Econometrica*, 83(6), 2293–2369. Retrieved August 22, 2025, from <https://www.jstor.org/stable/43866414>
- Gayle, G.-L., & Miller, R. A. (2009). Has Moral Hazard Become a More Important Factor in Managerial Compensation? *American Economic Review*, 99(5), 1740–1769. <https://doi.org/10.1257/aer.99.5.1740>
- Goshen, Z., Hamdani, A., & Lund, D. S. (2024). Fixing MFW: Fairness and Vision in Controller Self-Dealing. <https://doi.org/10.2139/ssrn.5061341>
- Grossman, S. J., & Stiglitz, J. E. (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, Vol. 70(No. 3), pp. 393–408.
- Hartzell, J. C., & Starks, L. T. (2003). Institutional Investors and Executive Compensation [Publisher: [American Finance Association, Wiley]]. *The Journal of Finance*, 58(6), 2351–2374. Retrieved August 21, 2025, from <https://www.jstor.org/stable/3648196>
- Holmstrom, B., & Milgrom, P. (1987). Aggregation and Linearity in the Provision of Intertemporal Incentives [Publisher: [Wiley, Econometric Society]]. *Econometrica*, 55(2), 303–328. <https://doi.org/10.2307/1913238>
- Hong, H., & Stein, J. C. (2007). Disagreement and the Stock Market [Publisher: American Economic Association]. *The Journal of Economic Perspectives*, 21(2), 109–128. Retrieved November 21, 2025, from <https://www.jstor.org/stable/30033720>
- Malmendier, U., & Tate, G. (2009). Superstar CEOs. *The Quarterly Journal of Economics*, 124(4), 1593–1638. <https://doi.org/10.1162/qjec.2009.124.4.1593>

- Miller, E. M. (1977). Risk, Uncertainty, and Divergence of Opinion [Publisher: [American Finance Association, Wiley]]. *The Journal of Finance*, 32(4), 1151–1168. <https://doi.org/10.2307/2326520>
- O'Brien, J. (2026). Wachtell, Lipton, Rosen & Katz Compensation Committee Guide - 2026.
- Shapley, L. S., & Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1, 111–130.
- Shleifer, A., & Vishny, R. W. (1997). The Limits of Arbitrage [Publisher: [American Finance Association, Wiley]]. *The Journal of Finance*, 52(1), 35–55. <https://doi.org/10.2307/2329555>
- Terviö, M. (2008). The Difference That CEOs Make: An Assignment Model Approach [Publisher: American Economic Association]. *The American Economic Review*, 98(3), 642–668. Retrieved August 19, 2025, from <https://www.jstor.org/stable/29730090>
- Thomas, R. S., Palmiter, A. R., & Cotter, J. F. (2012). Dodd-Frank's Say on Pay: Will It Lead to a Greater Role for Shareholders in Corporate Governance? *Cornell Law Review*, 97(5), 1213–1266. <https://scholarship.law.cornell.edu/clr/vol97/iss5/6>
- Yan, Y., Tong, Y., & Wang, Y. (2025). Validity of CARA function under expected utility. *Economics Letters*, 247, 112075. <https://doi.org/10.1016/j.econlet.2024.112075>

A Appendix: Holmstrom-Milgrom Framework and Derivation of Investor Adverse Selection Model

Firm's Problem

First, let's start with the production function for the firm. Firm's value is $V(a, \epsilon_{FIRM})$. The “effort” or “action” taken by the agent (CEO) is $a \in [0, \infty)$, with a lower value corresponding to “shirking” or low effort and a higher value corresponding to “high effort.” The uncertainty, ϵ_{FIRM} , is a mean-zero, finite variance noise term. The action, a , is taken first, then the uncertainty, ϵ_{FIRM} , is realized. Neither of $\{a, \epsilon_{FIRM}\}$ are observable to the principals (investors), though a is known by the agent (CEO).

While it is clear that $\frac{\partial V}{\partial a} > 0$, the exact functional form of $V(a, \epsilon)$ may change certain results (see Edmans and Gabaix (2016)).

The agent is paid a wage contingent on firm value, $c(V)$, which is subject to $0 \leq c(V) \leq V$. The agent's objective function is then:

$$\mathbb{E}[U] = \mathbb{E}[u(v(c) - g(a))]$$

where $u(\dots)$ is the agent's utility function, $v(c)$ is the agent's felicity function⁵⁶ and $g(a)$ is the agent's cost of effort. It is assumed that $g'(a) > 0$, $g''(a) > 0$.

Principal's Problem

It is conventional to assume that the principals (investors) are risk-neutral because they are well diversified. Hence, their objective is to maximize:

$$\begin{aligned} \max_{\{c(\cdot), a\}} \quad & \mathbb{E}[V(a) - c(V(a))] \\ \text{s.t.,} \quad & \mathbb{E}[u(v(c(V(\hat{a}))) - g(\hat{a}))] \geq \omega \\ & \& a \in \arg \max_{\hat{a}} \mathbb{E}[u(v(c(V(\hat{a}))) - g(\hat{a}))] \end{aligned} \tag{14}$$

Optimal Contract

From Lemma 2 in Appendix B, the solution to the agent's problem is:

⁵⁶I.e., the “ $v(d)$ ” notation is used to denote the agent's utility from cash.

$$a^* = \frac{\theta b^*}{g} \quad (15)$$

Plugging (15) into (14), setting the participation constraint to bind, evaluating the expectation on the left-hand side and equating the exponents yields an objective function of:

$$Z + \frac{\theta b^2}{g} - \omega - \frac{\theta^2 b^2}{2g} - \frac{\eta \theta^2 \sigma_{FIRM}^2}{2} \quad (16)$$

Taking the first order conditions of (16) with respect to θ gives:

$$\theta^* = \frac{1}{1 + g\eta \left(\frac{\sigma_{FIRM}}{b}\right)^2} \quad (17)$$

Revising the agent's participation constraint again with (17) yields:

$$\phi^* = \omega - \theta Z - \frac{\theta^2 b^2}{2g} + \frac{\eta}{2} \theta^2 \sigma_{FIRM}^2 \quad (18)$$

B Appendix: Proofs

This appendix collects the proofs supporting the model, in the order the corresponding results appear in the text: the CEO's effort choice (Section 3.2.4's contracting environment), the monotonicity of investor valuations underlying market clearing, single-peakedness of preferences over contracts and the median-voter result (Section 3.2.4), adverse selection of the median shareholder (Section 3.2.5), and the non-invertibility results (Sections 3.2.6–3.2.7). Statements proved here for the first time are given numbered environments; results stated in the main text are proved under their original numbering.

B.1 CEO Effort

Lemma 2 (CEO effort choice). *Under an implemented linear contract $\tilde{c}(a) = \tilde{\phi} + \tilde{\theta} V(a, \epsilon^{FIRM})$ whose participation constraint is satisfied, the CEO's optimal effort is $a^* = \tilde{\theta} b^* / g$.*

Proof. The CEO solves

$$\max_a \mathbb{E}[-\exp(-\eta [\tilde{c}(a) - \kappa(a)])], \quad (19)$$

where $\kappa(a) = \frac{1}{2} g a^2$ is the effort cost. Substituting the firm value equation, $V = Z + b^* a + \epsilon^{FIRM}$,

$$\tilde{c}(a) - \kappa(a) = \tilde{\phi} + \tilde{\theta} Z + \tilde{\theta} b^* a - \frac{1}{2} g a^2 + \tilde{\theta} \epsilon^{FIRM}.$$

Since $\epsilon^{FIRM} \sim \mathcal{N}(0, \sigma_{FIRM}^2)$ and utility is CARA, maximizing expected utility is equivalent to maximizing the

certainty equivalent

$$CE(a) = \tilde{\phi} + \tilde{\theta}Z + \tilde{\theta}b^*a - \frac{1}{2}ga^2 - \frac{\eta}{2}\tilde{\theta}^2\sigma_{FIRM}^2,$$

in which only the two middle terms involve a ; the risk adjustment does not, because effort shifts the mean of pay but not its variance. The first-order condition $\tilde{\theta}b^* - ga = 0$ is sufficient by strict concavity ($CE''(a) = -g < 0$) and yields $a^* = \tilde{\theta}b^*/g$. $\square$

B.2 Market Clearing: Valuations Are Increasing in Perceived Skill

Lemma 3 (Valuations are strictly increasing in perceived skill). *Let $\pi_0(\tilde{b}) \equiv \max_{\theta} [\theta\tilde{b}^2/g - \theta^2\tilde{b}^2/(2g) - (\eta/2)\theta^2\sigma_{FIRM}^2] - \omega$ denote the skill-dependent component of an investor's expected profit under the HM-optimal contract for perceived skill $\tilde{b}$. Then $\pi'_0(\tilde{b}) > 0$ for all $\tilde{b} > 0$.*

Proof. By the envelope theorem, $\pi'_0(\tilde{b})$ equals the partial derivative of the bracketed objective with respect to $\tilde{b}$, evaluated at the maximizing slope $\theta^*(\tilde{b}) = \tilde{b}^2/(\tilde{b}^2 + g\eta\sigma_{FIRM}^2)$:

$$\pi'_0(\tilde{b}) = \frac{2\theta^*\tilde{b}}{g} - \frac{(\theta^*)^2\tilde{b}}{g} = \frac{\theta^*\tilde{b}}{g} (2 - \theta^*) > 0,$$

since $\theta^* \in (0, 1)$ for all $\tilde{b} \in (0, \infty)$. Hence per-share valuations $w = [Z + \pi_0(\tilde{b})]/s$ are strictly increasing in $\tilde{b}$, and ranking investors by valuation is equivalent to ranking them by signal. $\square$

B.3 Contract Setting: Single-Peaked Preferences and the Median Voter

Lemma 4 (Single-peaked preferences over contracts). *For each investor i with perceived skill $\tilde{b}_i > 0$, the payoff $\hat{\Pi}_i(\tilde{b}_j)$ from implementing the HM-optimal contract of an investor with perceived skill $\tilde{b}_j$ is single-peaked in $\tilde{b}_j$ on $(0, \infty)$, with its unique peak at $\tilde{b}_j = \tilde{b}_i$.*

Proof. The proof proceeds in three steps: (i) I define the “cross-evaluation” profit function that captures investor i 's assessment of a contract designed for belief $\tilde{b}_j$; (ii) I show this function is globally concave in the contract slope θ ; and (iii) I use the monotonicity of the mapping $\tilde{b} \mapsto \theta^*(\tilde{b})$ to lift single-peakedness from θ -space to $\tilde{b}$ -space.

Step 1: Cross-evaluation profit. Suppose investor j proposes the HM-optimal contract for perceived skill $\tilde{b}_j$. This contract has slope:

$$\theta(\tilde{b}_j) = \frac{\tilde{b}_j^2}{\tilde{b}_j^2 + g\eta\sigma_{FIRM}^2} \tag{20}$$

Now consider investor i , who believes the CEO's skill is $\tilde{b}_i$. If investor j 's contract is implemented, investor i forecasts

(by Lemma 2, applied at their own belief) that the CEO will exert effort:⁵⁷

$$\tilde{a}_i(\theta_j) = \frac{\theta(\tilde{b}_j) \cdot \tilde{b}_i}{g} \quad (21)$$

Investor i 's expected profit from this contract is therefore:

$$\Pi_i(\tilde{b}_j) \equiv \mathbb{E}_i[V(\tilde{a}_i(\theta_j), \epsilon^{FIRM}) - c(\theta(\tilde{b}_j), V)] \quad (22)$$

Substituting the firm value equation (1), the linear contract $c = \phi + \theta V$, and evaluating the expectation (using $\mathbb{E}[\epsilon^{FIRM}] = 0$), we obtain:

$$\Pi_i(\theta_j) = \underbrace{Z + \frac{\theta_j \tilde{b}_i^2}{g}}_{\text{expected firm value}} - \underbrace{\left(\omega + \frac{\theta_j^2 \tilde{b}_i^2}{2g} + \frac{\eta \theta_j^2 \sigma_{FIRM}^2}{2} \right)}_{\text{expected compensation cost}} \quad (23)$$

where $\theta_j \equiv \theta(\tilde{b}_j)$ and I have used the binding participation constraint to substitute for ϕ .⁵⁸

Note the key distinction from the standard single-investor problem: the contract slope θ_j is determined by investor j 's belief $\tilde{b}_j$, while the effort forecast and value assessment use investor i 's belief $\tilde{b}_i$. When $\tilde{b}_j = \tilde{b}_i$, this reduces to investor i 's own optimal contract.

Step 2: Global concavity in θ . I now show that Π_i is globally concave in θ_j , treating $\tilde{b}_i$ as fixed. Differentiating equation (23) with respect to θ_j :

$$\frac{\partial \Pi_i}{\partial \theta_j} = \frac{\tilde{b}_i^2}{g} - \frac{\theta_j \tilde{b}_i^2}{g} - \eta \theta_j \sigma_{FIRM}^2 \quad (24)$$

The first term is the marginal revenue of increasing incentive pay (more effort $\rightarrow$ higher firm value), and the remaining terms are the marginal cost (higher compensation expense, comprising both the direct pay-for-performance cost and the risk premium). Setting this equal to zero yields investor i 's preferred contract slope:

$$\theta_i^* = \frac{\tilde{b}_i^2}{\tilde{b}_i^2 + g \eta \sigma_{FIRM}^2} \quad (25)$$

which is, as expected, the HM-optimal slope for perceived skill $\tilde{b}_i$. Differentiating equation (24) again:

$$\frac{\partial^2 \Pi_i}{\partial \theta_j^2} = -\frac{\tilde{b}_i^2}{g} - \eta \sigma_{FIRM}^2 < 0 \quad \forall \theta_j \quad (26)$$

Since the second derivative is strictly negative everywhere—not just at the optimum— Π_i is *globally* concave in θ_j .

This means Π_i has a unique maximum at θ_i^* and is strictly decreasing as θ_j moves away from θ_i^* in either direction.

⁵⁷The CEO's actual effort depends on the CEO's true skill b^* , not on any investor's belief. But in the contract-setting stage, investor i does not know b^* and forecasts effort using their own belief $\tilde{b}_i$. The single-peakedness result concerns investor preferences *at the time of voting*, which are formed under their own beliefs; see Assumption 1.

⁵⁸Writing θ_j in place of $\theta(\tilde{b}_j)$ is purely for notational convenience; θ_j remains a function of $\tilde{b}_j$ throughout.

Therefore, investor i 's preferences over contracts are single-peaked in θ .

Step 3: Lifting to $\tilde{b}$ -space. Single-peakedness in θ does not immediately imply single-peakedness in $\tilde{b}$, because the mapping $\tilde{b}_j \mapsto \theta(\tilde{b}_j)$ could in principle be non-monotone, creating multiple peaks. I now verify that this mapping is strictly increasing, which rules out this concern. Differentiating equation (20) with respect to $\tilde{b}_j$:

$$\frac{\partial \theta}{\partial \tilde{b}_j} = \frac{2g\eta\sigma_{FIRM}^2 \tilde{b}_j}{(\tilde{b}_j^2 + g\eta\sigma_{FIRM}^2)^2} > 0 \quad \forall \tilde{b}_j > 0 \quad (27)$$

Hence $\theta(\cdot)$ is a strictly increasing, continuous function on $\tilde{b}_j \in (0, \infty)$.

Define investor i 's payoff as a function of the policy variable $\tilde{b}_j$:

$$\hat{\Pi}_i(\tilde{b}_j) \equiv \Pi_i(\theta(\tilde{b}_j))$$

Since $\Pi_i(\theta)$ is single-peaked in θ with peak at θ_i^* , and $\theta(\cdot)$ is a strictly increasing bijection from $(0, \infty)$ to $(0, 1)$, the composition $\hat{\Pi}_i(\tilde{b}_j)$ is single-peaked in $\tilde{b}_j$ with peak at the unique $\tilde{b}_j^*$ satisfying $\theta(\tilde{b}_j^*) = \theta_i^*$ —that is, at $\tilde{b}_j^* = \tilde{b}_i$.

To see this explicitly: for any $\tilde{b}_j' > \tilde{b}_j > \tilde{b}_i$, we have $\theta(\tilde{b}_j') > \theta(\tilde{b}_j) > \theta_i^*$, and since Π_i is strictly decreasing in θ above θ_i^* , it follows that $\hat{\Pi}_i(\tilde{b}_j') < \hat{\Pi}_i(\tilde{b}_j) < \hat{\Pi}_i(\tilde{b}_i)$. An analogous argument holds for $\tilde{b}_j < \tilde{b}_i$. Hence $\hat{\Pi}_i$ is strictly increasing for $\tilde{b}_j < \tilde{b}_i$ and strictly decreasing for $\tilde{b}_j > \tilde{b}_i$, which is the definition of single-peakedness. $\square$

Remark 1 (Application of the median voter theorem). *Since shareholders' preferences over the policy variable $\tilde{b}_j$ are single-peaked (Lemma 4), the median voter theorem (Black, 1948) applies directly: the contract that wins a majority against any alternative is the one corresponding to the median shareholder's belief, $\tilde{b}_{\bar{i}}$, where $\bar{i}$ denotes the shareholder with the median signal among the s shareholders who purchased shares in the market-clearing stage.*

B.4 Adverse Selection of the Median Shareholder

Proposition 2 (Adverse selection of the median shareholder). *Let $N > s \geq 2$ and let signals $\{\tilde{b}_i\}_{i=1}^N$ be drawn i.i.d. from $N(b^*, \sigma_{INV}^2)$. After market clearing, the s shareholders are the investors with the s highest signals. Let $\bar{i}$ denote the median shareholder, i.e., the shareholder with the $\lceil s/2 \rceil$ -th highest signal among the s shareholders (equivalently, the investor with the $\lceil s/2 \rceil$ -th highest signal in the full population). Then:*

- (i) $P(\tilde{b}_{\bar{i}} > b^*) > \frac{1}{2}$; and
- (ii) $\mathbb{E}[\tilde{b}_{\bar{i}}] > b^*$.

That is, the median shareholder's signal is both more likely to exceed the true CEO skill than to fall below it, and biased upward in expectation.

Proof. Let $k \equiv \lceil s/2 \rceil$ denote the rank (from the top) of the median shareholder in the full population. Since each $\tilde{b}_i = b^\star + \epsilon_i^{INV}$ with $\epsilon_i^{INV} \sim \mathcal{N}(0, \sigma_{INV}^2)$, we can write:

$$\tilde{b}_{(k:N)} = b^\star + \sigma_{INV} \cdot Z_{(k:N)}$$

where $Z_{(k:N)}$ is the k -th descending order statistic of N i.i.d. standard normal draws. Both claims therefore reduce to properties of $Z_{(k:N)}$.

Part (i). I want to show $P(Z_{(k:N)} > 0) > 1/2$, or equivalently, that $P(\text{at least } k \text{ of the } N \text{ draws exceed } 0) > 1/2$.

Let $M = \sum_{i=1}^N \mathbb{1}\{Z_i > 0\}$ denote the number of standard normal draws that exceed zero. Since $P(Z_i > 0) = 1/2$ for each draw, $M \sim \text{Binomial}(N, 1/2)$. The event $\{Z_{(k:N)} > 0\}$ is equivalent to $\{M \geq k\}$, so:

$$P(Z_{(k:N)} > 0) = P(M \geq k) = \sum_{m=k}^N \binom{N}{m} \left(\frac{1}{2}\right)^N$$

We need $P(M \geq k) > 1/2$, which is equivalent to $P(M \leq k-1) < 1/2$. Since $M \sim \text{Binomial}(N, 1/2)$, the distribution of M is symmetric around $N/2$. Therefore $P(M \geq k) > 1/2$ whenever $k \leq \lfloor N/2 \rfloor$.⁵⁹

It remains to verify that $k = \lceil s/2 \rceil \leq \lfloor N/2 \rfloor$. Since $s < N$:

$$k = \lceil s/2 \rceil \leq \lceil (N-1)/2 \rceil = \lfloor N/2 \rfloor$$

where the last equality uses a standard identity for integers.⁶⁰ Hence $P(\tilde{b}_i > b^\star) > 1/2$. This establishes part (i).

Part (ii). I want to show $\mathbb{E}[Z_{(k:N)}] > 0$. By the symmetry of the standard normal distribution ($\phi(z) = \phi(-z)$ and $\Phi(-z) = 1 - \Phi(z)$), the order statistics satisfy:

$$Z_{(k:N)} \stackrel{d}{=} -Z_{(N-k+1:N)} \tag{28}$$

That is, the k -th from the top has the same distribution as the negative of the k -th from the bottom, so $\mathbb{E}[Z_{(k:N)}] = -\mathbb{E}[Z_{(N-k+1:N)}]$. Expected order statistics are strictly decreasing in rank for any continuous distribution, $\mathbb{E}[Z_{(1:N)}] > \mathbb{E}[Z_{(2:N)}] > \dots > \mathbb{E}[Z_{(N:N)}]$.⁶¹ From part (i), $k \leq \lfloor N/2 \rfloor < (N+1)/2$, so $k < N - k + 1$ and therefore

$$\mathbb{E}[Z_{(k:N)}] > \mathbb{E}[Z_{(N-k+1:N)}] = -\mathbb{E}[Z_{(k:N)}],$$

⁵⁹For even N , the binomial($N, 1/2$) distribution places positive mass at $N/2$, so $P(M \geq N/2) > 1/2$ and the condition is $k \leq N/2$. For odd N , $P(M \geq (N+1)/2) = 1/2$ by symmetry, and the condition is $k \leq (N-1)/2$. In both cases, $k \leq \lfloor N/2 \rfloor$ suffices.

⁶⁰For even N : $\lceil (N-1)/2 \rceil = N/2 = \lfloor N/2 \rfloor$. For odd N : $\lceil (N-1)/2 \rceil = (N-1)/2 = \lfloor N/2 \rfloor$.

⁶¹See, e.g., David and Nagaraja (2003), Chapter 4.

which forces $\mathbb{E}[Z_{(k:N)}] > 0$. Hence

$$\mathbb{E}[\tilde{b}_{\bar{i}}] = b^* + \sigma_{INV} \cdot \mathbb{E}[Z_{(k:N)}] > b^*.$$

□

Corollary 1 (Adverse selection distorts effort and reduces expected profit). *Under the conditions of Proposition 2 (in particular $N > s \geq 2$ and $\sigma_{INV} > 0$) and Assumption 1, the compensation contract set by the median shareholder $\bar{i}$:*

(i) *provides strictly excessive incentive pay more likely than not: $P(\theta(\tilde{b}_{\bar{i}}) > \theta(b^*)) > \frac{1}{2}$;*

(ii) *induces CEO effort strictly above the correct-perception benchmark more likely than not: $P(a^*(\tilde{b}_{\bar{i}}) > a^*(b^*)) > \frac{1}{2}$; and*

(iii) *yields expected firm profit strictly below the correct-perception benchmark: $\mathbb{E}[\pi(\tilde{b}_{\bar{i}}, b^*)] < \pi(b^*, b^*)$.*

Proof. Part (i). By equation (27), $\theta(\cdot)$ is strictly increasing on $(0, \infty)$. Since $b^* > 0$, on the event $\{\tilde{b}_{\bar{i}} > b^*\}$ we have $\theta(\tilde{b}_{\bar{i}}) > \theta(b^*)$, so $\{\tilde{b}_{\bar{i}} > b^*\} \subseteq \{\theta(\tilde{b}_{\bar{i}}) > \theta(b^*)\}$ and therefore

$$P(\theta(\tilde{b}_{\bar{i}}) > \theta(b^*)) \geq P(\tilde{b}_{\bar{i}} > b^*) > \frac{1}{2},$$

where the final inequality is Proposition 2(i). (The containment, rather than equality, of events is all that is required; realizations with $\tilde{b}_{\bar{i}} \leq b^*$ can only add to the left-hand event, never subtract from it.) This establishes part (i).

Part (ii). By Lemma 2, the CEO's effort under contract slope $\theta(\tilde{b}_{\bar{i}})$ and true skill b^* is $a^*(\tilde{b}_{\bar{i}}) = \theta(\tilde{b}_{\bar{i}}) b^*/g$, a strictly increasing linear transformation of $\theta(\tilde{b}_{\bar{i}})$. Hence $\{a^*(\tilde{b}_{\bar{i}}) > a^*(b^*)\} = \{\theta(\tilde{b}_{\bar{i}}) > \theta(b^*)\}$, and the claim follows from part (i).

Part (iii). Under Assumption 1, base pay binds the CEO's participation constraint at the true productivity, so realized expected profit at contract slope $t \in [0, 1]$ is

$$\Pi(t) = Z + \frac{t(b^*)^2}{g} - \omega - \frac{t^2(b^*)^2}{2g} - \frac{\eta t^2 \sigma_{FIRM}^2}{2}.$$

Π is strictly concave, $\Pi''(t) = -(b^*)^2/g - \eta \sigma_{FIRM}^2 < 0$, and its first-order condition $\Pi'(t) = (b^*)^2/g - t(b^*)^2/g - \eta t \sigma_{FIRM}^2 = 0$ is solved uniquely by $t^* = (b^*)^2 / ((b^*)^2 + g \eta \sigma_{FIRM}^2) = \theta(b^*)$. Hence t^* is the unique global maximizer of Π , and

$$\Pi(t) < \Pi(t^*) \quad \text{for every } t \neq t^*. \quad (29)$$

Let $T \equiv \theta(\tilde{b}_{\bar{i}})$. Since $\theta(b) = b^2/(b^2 + g \eta \sigma_{FIRM}^2)$ is strictly increasing in b^2 , we have $T = t^*$ if and only if $\tilde{b}_{\bar{i}}^2 = (b^*)^2$, i.e., $\tilde{b}_{\bar{i}} \in \{-b^*, b^*\}$. The median shareholder's signal $\tilde{b}_{\bar{i}} = b^* + \sigma_{INV} Z_{(k:N)}$ is absolutely continuous (an order statistic of i.i.d. normal draws has a density), so $P(\tilde{b}_{\bar{i}} \in \{-b^*, b^*\}) = 0$ and thus $P(T \neq t^*) = 1$. Because $T \in (0, 1)$ and Π is continuous on the compact interval $[0, 1]$, $\mathbb{E}[\Pi(T)] < \infty$. Combining with (29), $\Pi(T) < \Pi(t^*)$ almost surely, and

taking expectations of an almost-surely strict inequality between an integrable random variable and a constant yields

$$\mathbb{E}[\pi(\tilde{b}_{\bar{t}}, b^*)] = \mathbb{E}[\Pi(T)] < \Pi(t^*) = \pi(b^*, b^*). \quad \square$$

Remark 2 (Expectation versions of (i)–(ii)). *The mean of $\theta(\tilde{b}_{\bar{t}})$ also exceeds $\theta(b^*)$ throughout the empirically relevant parameter region: numerical evaluation against the exact order-statistic density finds $\mathbb{E}[\theta(\tilde{b}_{\bar{t}})] > \theta(b^*)$ whenever $\sigma_{INV} \leq b^*$, and whenever $\theta(b^*) \leq 1/4$ (the convex region of θ) even for σ_{INV} as large as $3b^*$. At the contract slopes observed in the estimation sample ($\hat{\theta}_{it} \sim 10^{-3}$ – 10^{-2}), the latter condition holds by three orders of magnitude. The expectation ordering can reverse only when extreme signal noise ($\sigma_{INV} \gtrsim 1.5b^*$) coincides with near-saturated contracts ($\theta(b^*) \gtrsim 0.94$) and weak selection ($s \approx N$)—and even there the median ordering of parts (i)–(ii) is unaffected.*

B.5 Non-Invertibility of the Market-Clearing Price

The results of Sections 3.2.6–3.2.7 are proved here under their original numbering. (The proof of Proposition 1 appears in the main text.)

Proof of Lemma 1. (i) The log-likelihood is built from $\ln \Gamma(\cdot)$, $\ln \Phi$, $\ln(1 - \Phi)$ and $\ln \phi$ evaluated at finite arguments; each is finite, so the likelihood is strictly positive for every admissible (b^*, ξ) (cf. the proof of Proposition 1). (ii) Define $F(\Delta, \xi) \equiv \ell(b_1^* - \Delta, \xi) - \ell(b_1^*, \xi_1)$. Then F is C^∞ on an open neighborhood of $(0, \xi_1)$, $F(0, \xi_1) = 0$, and by hypothesis $\partial F / \partial \xi|_{(0, \xi_1)} \neq 0$. The Implicit Function Theorem yields an open interval $(-\bar{\Delta}, \bar{\Delta})$ and a unique C^1 function $\xi^*(\Delta)$ with $\xi^*(0) = \xi_1$ and $F(\Delta, \xi^*(\Delta)) = 0$ throughout. Setting $\mathcal{U} = (b_1^* - \bar{\Delta}, b_1^* + \bar{\Delta})$ gives the result; symmetry of the interval around $\Delta = 0$ delivers alternatives on both sides of the truth. $\square$

Because market thickness is the paper’s leading instance, the construction underlying Corollary 2 is next carried out explicitly, including the treatment of integer-valued thickness.

Corollary 2 (Unknown market thickness). *Suppose the investor knows (s, σ_{INV}, Z) but not the number of potential investors. Then the conclusions of Lemma 1 hold with $\xi = \nu$: every skill level in an open neighborhood of the truth, on either side, is exactly as likely as the truth once market thickness is adjusted. The relevant derivative is*

$$\frac{\partial \ell}{\partial \nu} = \psi(\nu + 1) - \psi(\nu - s + 1) + \ln \Phi(z_1),$$

where $\psi(\cdot) = \Gamma'(\cdot)/\Gamma(\cdot)$ is the digamma function. The digamma difference is strictly positive (as ψ is strictly increasing on $(0, \infty)$ and $s \geq 1$) and $\ln \Phi(z_1)$ is strictly negative, so the derivative is nonzero outside a measure-zero set of (z_1, ν, s) configurations.

Proof of Corollary 2 (explicit construction). The proof applies the Implicit Function Theorem to the log-likelihood ratio, treated as a smooth function of CEO skill and market thickness. I proceed in three steps.

Step 1: The log-likelihood ratio as a smooth function with a known zero. Define $z(\Delta) \equiv (\hat{b} - b_1^* + \Delta)/\sigma_{INV}$, so that $z(0) = z_1 \equiv (\hat{b} - b_1^*)/\sigma_{INV}$ corresponds to the baseline scenario and $z(\Delta)$ for $\Delta \neq 0$ corresponds to an alternative scenario with $b_2^* = b_1^* - \Delta$. To allow the number of investors to vary continuously, extend the factorial function to positive reals via the Gamma function.

The log-likelihood ratio is:

$$\begin{aligned}
F(\Delta, \nu) \equiv \ln \left[\frac{f_{(s:\nu)}(\hat{b} \mid b_1^* - \Delta, \nu)}{f_{(s:N_1)}(\hat{b} \mid b_1^*, N_1)} \right] &= \ln \left[\frac{\Gamma(\nu + 1) / \Gamma(\nu - s + 1)}{\Gamma(N_1 + 1) / \Gamma(N_1 - s + 1)} \right] \\
&+ (\nu - s) \ln \Phi(z(\Delta)) - (N_1 - s) \ln \Phi(z_1) \\
&+ (s - 1) [\ln(1 - \Phi(z(\Delta))) - \ln(1 - \Phi(z_1))] \\
&+ \ln \phi(z(\Delta)) - \ln \phi(z_1)
\end{aligned} \tag{30}$$

where $\nu \in (s, \infty)$ is a real-valued stand-in for N_2 . All component functions— $\ln \Gamma$, $\ln \Phi$, $\ln(1 - \Phi)$, and $\ln \phi$ —are smooth on their respective domains, so F is C^∞ on $\mathbb{R} \times (s, \infty)$.⁶²

At the anchor point $(\Delta, \nu) = (0, N_1)$, every term cancels:

$$F(0, N_1) = 0$$

This is simply the observation that, if we do not change CEO skill or the number of investors, the two scenarios are identical.

Step 2: The Implicit Function Theorem. We wish to solve $F(\Delta, \nu) = 0$ for ν as a function of Δ in a neighborhood of the anchor point $(0, N_1)$. By the Implicit Function Theorem, a sufficient condition is that $\partial F / \partial \nu \neq 0$ at this point.

Differentiating (30) with respect to ν :

$$\frac{\partial F}{\partial \nu} = \psi(\nu + 1) - \psi(\nu - s + 1) + \ln \Phi(z(\Delta)) \tag{31}$$

where $\psi(\cdot) = \Gamma'(\cdot)/\Gamma(\cdot)$ is the digamma function.⁶³ At the anchor point:

$$\left. \frac{\partial F}{\partial \nu} \right|_{(0, N_1)} = \underbrace{\psi(N_1 + 1) - \psi(N_1 - s + 1)}_{> 0} + \underbrace{\ln \Phi(z_1)}_{< 0} \tag{32}$$

The digamma difference equals $\sum_{j=1}^s (N_1 - s + j)^{-1} > 0$.⁶⁴ The second term is strictly negative since $\Phi(z_1) \in (0, 1)$.

These two terms have opposite signs, and their sum is generically nonzero.⁶⁵

⁶²Smoothness of $\ln \Phi(z)$ and $\ln(1 - \Phi(z))$ requires $\Phi(z) \in (0, 1)$, which holds for all finite z , and $\ln \phi(z) = -\frac{1}{2} \ln(2\pi) - z^2/2$ is smooth everywhere.

⁶³This uses $\frac{d}{d\nu} \ln[\Gamma(\nu + 1)/\Gamma(\nu - s + 1)] = \psi(\nu + 1) - \psi(\nu - s + 1)$, which follows from the chain rule applied to $\ln \Gamma$.

⁶⁴This follows from the identity $\psi(n + 1) - \psi(m + 1) = \sum_{j=m+1}^n j^{-1}$ for positive integers $n > m$.

⁶⁵The derivative vanishes only on a measure-zero set of (z_1, N_1, s) configurations. In the non-generic case, one can perturb the anchor to a nearby

Assuming $\partial F / \partial v|_{(0, N_1)} \neq 0$, the Implicit Function Theorem guarantees the existence of:

- an open interval $(-\bar{\Delta}, \bar{\Delta})$ around $\Delta = 0$, and
- a unique C^1 function $v^* : (-\bar{\Delta}, \bar{\Delta}) \rightarrow \mathbb{R}$

such that $F(\Delta, v^*(\Delta)) = 0$ for all $\Delta \in (-\bar{\Delta}, \bar{\Delta})$, with $v^*(0) = N_1$. Setting $\mathcal{U} = (b_1^* - \bar{\Delta}, b_1^* + \bar{\Delta})$ and $\xi^*(b_2^*) = v^*(b_1^* - b_2^*)$ establishes the Corollary's claim under the Gamma-extension: every $b_2^* \in \mathcal{U}$, on either side of b_1^* , is exactly as likely as the truth at the adjusted (real-valued) thickness. Alternatives with $b_2^* < b_1^*$ (overvaluation) and $b_2^* > b_1^*$ (undervaluation) are both indistinguishable from the truth.

Step 3: Integer-valued thickness. When the number of investors is restricted to integers, two versions of the result survive.

(a) *Within any tolerance, everywhere nearby.* Set $N_2 = \lfloor v^*(\Delta) \rfloor$, the nearest integer to $v^*(\Delta)$. Since $v^*(\Delta) \rightarrow N_1$ as $\Delta \rightarrow 0$, for sufficiently small $|\Delta|$ the nearest integer is N_1 itself, and continuity of F gives $F(\Delta, N_2) \rightarrow F(0, N_1) = 0$. Hence for every tolerance $\varepsilon > 0$ there exist alternatives b_2^* on both sides of b_1^* whose likelihood ratio against the truth lies within ε of one at integer thickness.

(b) *Exact, at integer crossings.* Along the implicit path,

$$\frac{dv^*}{d\Delta} = -\frac{\partial F / \partial \Delta}{\partial F / \partial v}, \quad \frac{\partial F}{\partial \Delta} = \frac{1}{\sigma_{INV}} \left[(\nu - s) h_+(z(\Delta)) - (s - 1) h_-(z(\Delta)) - z(\Delta) \right],$$

where $h_+(z) = \phi(z) / \Phi(z)$ and $h_-(z) = \phi(z) / (1 - \Phi(z))$. The numerator is generically nonzero (it vanishes only on a measure-zero set of (z_1, N_1, s) configurations), so $v^*(\Delta)$ is generically strictly monotone in Δ and sweeps an interval of real thickness values through N_1 . Wherever the continued path crosses an integer value N_2 , the corresponding alternative $(b_1^* - \Delta, N_2)$ satisfies the likelihood-ratio equality *exactly*.

In either form, the observed price provides no discriminating power against alternatives on either side of the truth, which is the content of Corollary 2. $\square$

C Appendix: Data

Table 8 summarizes the data sources used in this paper.

D Appendix: Empirics

D.1 Say-on-Pay Vote Data Construction

Say-on-pay vote tabulations are disclosed in Item 5.07 of Form 8-K. I retrieve, from SEC EDGAR, every 8-K (and amendment) reporting Item 5.07 results filed by the estimation-sample firms over 2010–2026—23,061 filings—and point where the derivative is nonzero.

Table 8: Data Sources and Uses

Source	Variables	Sample Period	Used In
S&P ExecuComp (via WRDS)	CEO total compensation (<i>tdc1</i>), variable compensation components, CEO tenure, firm identifiers	2000–2025	Reduced-form regressions (§4.1); structural estimation (§5)
CRSP	Monthly stock returns, market capitalization, share price, trading volume, bid-ask spreads	2000–2025	Structural estimation (§5); asset pricing tests (§6.2)
FactSet (via WRDS)	Institutional ownership by investor type (active, passive, index)	2013–2022	Reduced-form regressions (§4.1)
Kenneth French Data Library	Fama–French factors (MKT, SMB, HML, RMW, CMA) and momentum (UMD); risk-free rate	2002–2024	Asset pricing tests (§6.2)
SEC EDGAR (Form 8-K, Item 5.07)	Say-on-pay vote tabulations (For/Against/Abstain)	2010–2026	Say-on-pay analysis (§6.1); construction in §D.1

Notes: All WRDS data were accessed in 2025. Firms with fewer than four CEO-year observations are excluded from the structural estimation. The reduced-form regression sample is restricted to 2013–2022, beginning with the first full year of comprehensive online Section 13 ownership filings.

extract the For, Against, and Abstain counts on the management say-on-pay proposal. Extraction is performed twice, independently: by a rule-based parser built around the recurring tabular and narrative formats of Item 5.07, with guards against the principal contamination risks (director-election tallies, shareholder proposals, and say-on-frequency votes, each of which places vote counts adjacent to the say-on-pay item); and by a large language model applied to the same filing text under a prompt requiring verbatim number extraction. The two methods agree exactly on both vote counts in 76.1% of filings where both produce an extraction; all disagreements, all near-failing votes, and a sample of formatting edge cases were manually adjudicated against the underlying filings. The final dataset contains 19,456 firm-year votes.

Two external checks support the extraction. First, the incidence of failed votes (1.9% below majority support) and of votes below the 70% proxy-advisor responsiveness threshold (6.3%) matches rates reported for the say-on-pay universe. Second, the distribution of approval shares reproduces the familiar heavy concentration near unanimity (median approval 95.5%), with failures occurring in every sample year. Replication code for the scrape, both extractors, the cross-validation, and the final assembly is included in the replication package.

D.1.1 Additional Say-on-Pay Analysis

Table 9 reports the full decomposition of the voting response underlying column (5) of Table 4, together with electorate-composition interactions. Three results sharpen the main-text finding. First, the signed orthogonal component moves votes nowhere ($t = -0.25$; $t = -1.38$ in the ownership subsample). Under symmetric single-peaked preferences, a linear term should be near zero; this is consistent with the model’s shape. More directly, it shows that votes carry no directional information about the belief error, which is the information that cleansing-by-approval would require. Second, the magnitude response in the benchmark component is a between-firm phenomenon. Within

a firm—the margin at which a voter would need to detect a given year’s distortion—the coefficient falls by two-thirds and loses conventional significance (-0.006 , $t = -1.9$). When the contract slope is purged flexibly rather than linearly (using within-year cubic and decile-dummy residualizations), the within-firm response disappears entirely ($t = -1.3$ and -0.8), while the between-firm response survives every purge ($t = -5.2$ and -3.5 , against -5.6 under the linear baseline; replication Table T26). Votes appear to incorporate the belief error only as accumulated performance slowly reveals it, not in the year the contract is ratified. Third, electorate composition does not rescue detection: the interaction of the orthogonal component with active institutional share is null ($t = 0.62$). The main-text results are therefore not an artifact of proxy-advisor-guided passive voting—though active institutions do dissent more per unit of observable misalignment ($t = -3.6$), consistent with voting governed by centralized stewardship policies rather than by the portfolio beliefs that drive holdings.

Table 9: Votes respond to observable contract terms, not the misperception component

	(1)	(2)	(3)	(4)	(5)	(6)
Orthogonalized $\tilde{\varepsilon}^\perp$ (SD)	-0.0006 (-0.25)					-0.0031 (-1.38)
$ \tilde{\varepsilon}^\perp $ (SD)		-0.0185*** (-5.65)	-0.0059* (-1.91)			
Misperception $\tilde{\varepsilon}$ (SD)				-0.0090*** (-3.14)	-0.0118*** (-4.45)	
Active institutional share				-0.0113 (-0.72)	-0.0076 (-0.48)	-0.0041 (-0.25)
$\tilde{\varepsilon} \times$ active share					-0.0267*** (-3.58)	
$\tilde{\varepsilon}^\perp \times$ active share						0.0065 (0.62)
Controls + Year FE	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓		✓	✓	✓
Firm FE			✓			
Observations	13,077	13,077	13,077	7,958	7,958	7,958
R^2	0.063	0.079	0.349	0.067	0.069	0.064

Notes: Dependent variable: say-on-pay approval share. $\tilde{\varepsilon}^\perp$ is misperception orthogonalized to the contract slope $\hat{\theta}$ within year (the component not observable from contract steepness); $|\cdot|$ denotes absolute value in SD units. Columns (2)–(3) control for absolute firm value growth. Columns (4)–(6) merge FactSet institutional ownership (2014–2023); active share = total institutional minus top-three passive, demeaned. t -statistics in parentheses, clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The board’s response channel also fails to transmit the belief error. Table 10 examines pay changes around low votes. The advisory mechanism does bind on observables: a below-70% vote is followed by pay growth 7.7 percentage points lower over the surrounding two years ($t = -3.0$), with the correction of the contract slope concentrated where the slope was observably extreme. But the responsiveness is unrelated to the misperception component once the observable slope is accounted for. The apparatus disciplines the terms of pay, but the belief distortion those terms embed passes through untouched.

Table 10: Board responses to low say-on-pay votes

	$\Delta \ln \text{ pay}$	$\Delta \ln \text{ pay}$	$\Delta \hat{\theta}$	$\Delta \hat{\theta}$
Low vote (< 70%)	-0.1169*** (-3.83)	-0.0769*** (-3.01)	0.0384 (0.92)	0.0525 (1.19)
Misperception $\tilde{\varepsilon}$ (SD)		-0.3071*** (-15.67)	-0.2516*** (-7.82)	-0.3144*** (-5.33)
Low vote $\times \tilde{\varepsilon}$		-0.0620 (-1.43)	-0.2737*** (-2.70)	0.0979 (0.86)
Low vote $\times \hat{\theta}_{pre}$ (SD)				-0.3894** (-2.43)
Controls + Year/Industry FE	✓	✓	✓	✓
Observations	11,240	11,240	11,240	11,240
R^2	0.205	0.297	0.371	0.380

Notes: Dependent variable: change in log total CEO compensation (columns (1)–(2)) or in the contract slope $\hat{\theta}$, SD units (columns (3)–(4)), from year $t - 1$ to $t + 1$ around a say-on-pay vote in year t . Low vote indicates approval below 70%. All columns control for the pre-vote slope level. t -statistics in parentheses, clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.1.2 Kink Scan Horse Race

Kink-location scan. The scan treats the location of the approval peak as a parameter to be estimated rather than imposed. Approval is first residualized on the controls in Table 4 (excess pay, log market capitalization, prior-year firm value growth, and year and industry fixed effects). The sorting variable is standardized to unit variance but not demeaned, so that zero remains the model’s zero-misperception point. For each candidate kink location c on a grid of 121 points spanning the 2nd to 98th percentiles of the standardized sorting variable (approximately $[-1.1, +3.1]$), I fit the two-slope specification—an intercept, $\max(x-c, 0)$, and $\min(x-c, 0)$ —by least squares and record the residual sum of squares. The fit-maximizing location is $\hat{c} = +0.10$ standard deviations. The competing prediction is located by the sample mean of the standardized sorting variable, $+0.34$ SD: if voters simply penalized contracts for being atypical, the approval peak would sit at the center of the distribution rather than at the model’s zero. The two locations are separated because the distribution is tilted— $\mu_{INV} > 0$ places the typical contract to the right of the aligned contract—so the adverse-selection finding is what makes the two hypotheses distinguishable.

Bootstrap. Inference on $\hat{c}$ comes from a firm-clustered bootstrap with $B = 1,000$ replications: firms are resampled with replacement, the residualization is re-estimated within each resampled panel, and the scan is re-run on the fixed grid, yielding one $\hat{c}$ per replication. The percentile 95% confidence interval is $[-0.31, +0.24]$, which contains the model’s zero. The exclusion of the firm-typical alternative is stronger than the interval suggests: the fit-maximizing kink falls below $+0.34$ SD in every one of the 1,000 replications (bootstrap $p < 0.001$). Two limitations bound the exercise. First, the grid resolution is approximately 0.03 SD, so $\hat{c}$ is grid-snapped. Second, the sorting variable is held at its point-estimate values within each replication, so the interval reflects sampling variation in votes and firms but not estimation uncertainty in the structural layer.

Benchmark horserace. While the scan locates the peak, the horserace asks whether the model’s distance measure carries information beyond simple alternatives a skeptic might propose. I construct four distances in standard-deviation units: the model’s $|\tilde{\epsilon}|$; the distance of the contract slope from the firm’s own historical mean slope; its distance from the industry-year median slope; and the absolute change in log total pay—the classic dissent trigger in the say-on-pay literature. The slope-based alternatives are highly correlated with the model’s measure (0.74 and 0.77 respectively, and 0.87 with each other); the pay-change trigger is nearly orthogonal to all three ($\rho \approx 0.2$). Each distance enters the approval regression alone and then all four jointly, using the full control set and firm-clustered standard errors:

	Model $ \tilde{\epsilon} $	Firm-history	Industry-median	$ \Delta \log \text{pay} $
Alone	−0.0194 (−5.49)	−0.0122 (−2.87)	−0.0097 (−2.83)	−0.0165 (−7.86)
Joint	−0.0171 (−4.55)	−0.0072 (−1.71)	+0.0096 (+2.42)	−0.0131 (−6.24)

Discussion. Three features of the joint estimates are key. First, the diagnostic of a measure carrying real signal is stability: the model’s coefficient barely moves from alone to joint (−0.0194 to −0.0171) despite racing against measures correlated at 0.74–0.77. In contrast, the firm-history benchmark becomes marginal and the industry-median benchmark flips sign entirely. The sign flip should not be over-read as economics—with three collinear distances in one regression, such suppression effects are typical of misspecified proxies—but that is the point: the naive benchmarks behave like noisy projections of the model’s measure, while the model’s measure behaves like the underlying object. Second, the pay-change trigger survives jointly ($t = -6.2$). This is not a defeat for the model; pay changes measure a different margin of dissent (voters react to levels moving, a channel long documented in the say-on-pay literature) and are nearly uncorrelated with all three alignment distances. The horserace claim is that the model’s measure dominates alternative measures of misalignment, not that misalignment is the only thing voters react to. Third, the scan and the horserace answer complementary objections: the scan shows the approval peak sits at the model’s zero rather than at the typical contract; the horserace shows that the curvature around that peak is organized by the fundamentals-benchmarked distance rather than by habit or peer comparison. Together, they support reading Figure 3 as evidence about alignment with estimated skill, not generic aversion to unusual pay. Implementation and archived replications: `sop_validation_tests.py` and `run_kink_bootstrap.py` in the replication package (outputs `sop_validation_horserace.csv`, `sop_validation_kinkscan.csv`, and Table T27).

D.2 Distinguishing Adverse Selection from Managerial Power

A natural objection to reading $\hat{\mu}_{INV} > 0$ as adverse selection is observational equivalence with managerial power: perhaps steep contracts and the dissent they attract reflect entrenched CEOs extracting rents rather than optimistic shareholder bases. Two features of the evidence already cut against this reading—dissent is two-sided around the

model's zero, where a rent-extraction account predicts dissent only against overpayment, and optimistically held firms suffer no return penalty, where rent extraction predicts one. This appendix adds a direct horse race. Table 11 augments the say-on-pay distance specification with six standard managerial-power proxies: the CEO pay slice (Bebchuk et al. (2011a)), CEO tenure, CEO–chair duality, CEO equity ownership, a dual-class indicator, and board co-option (Coles et al. (2014), whose posted data extend through 2022). The proxies carry dissent of their own—approval falls with the pay slice, tenure, duality, and co-option—so shareholder votes do police managerial power. But they police it *separately*: the misalignment-distance coefficient is essentially unchanged by any proxy individually and moves by three percent when all six enter jointly on a common sample (-0.0311 to -0.0301 , both $t < -5.5$). Reverse regressions sharpen the point: the six proxies together explain at most 2.2 percentage points of the variance in estimated misperception beyond size and year effects. Managerial power is real and visible in these data; it is also nearly orthogonal to what the structural estimation measures.

Table 11: Misalignment dissent survives managerial-power proxies

Specification	$ \bar{\varepsilon} $ coef.	(t)	Proxy coef.	(t)
Baseline (full vote sample)	-0.0192***	(-5.52)		
+ CEO pay slice	-0.0188***	(-5.43)	-0.0107***	(-4.71)
+ CEO tenure	-0.0186***	(-5.31)	-0.0074***	(-3.96)
+ CEO–chair duality	-0.0189***	(-5.48)	-0.0051***	(-2.82)
+ CEO ownership %	-0.0189***	(-5.47)	-0.0008	(-0.47)
+ dual-class	-0.0193***	(-5.64)	0.0019	(1.08)
+ board co-option	-0.0301***	(-5.66)	-0.0083***	(-5.03)
Baseline (co-option sample)	-0.0311***	(-5.82)		
All six proxies jointly	-0.0301***	(-5.52)		
<i>Reverse regressions (misperception on all six proxies + size, year FE):</i>				
signed $\hat{\varepsilon}$: R^2 0.190 $\rightarrow$ 0.212 ($\Delta R^2 = 0.022$)				$n = 8, 104$
distance $\hat{\varepsilon}$: R^2 0.223 $\rightarrow$ 0.235 ($\Delta R^2 = 0.012$)				$n = 8, 104$

Notes: Rows add one standardized managerial-power proxy at a time to the say-on-pay distance specification (Table 4, column (4)): CEO pay slice (CEO share of top-five executive pay), CEO tenure, CEO–chair duality (from ExecuComp titles), CEO equity ownership (% of shares), a dual-class indicator (multiple CRSP share classes), and board co-option (fraction of directors appointed after the sitting CEO took office, from the Coles–Daniel–Naveen data, available through 2022). The co-option rows use the smaller merged sample; the same-sample baseline is shown for comparison. Reverse regressions report how much of the variance in estimated misperception the six proxies jointly explain beyond size and year effects. t -statistics clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.3 Competing Explanations: Predictions and Evidence

Table 12 collects the distinguishing predictions of the four candidate explanations for $\hat{\mu}_{INV} > 0$ and the corresponding evidence.

D.4 Implementation of the Policy Counterfactuals

This appendix describes the computation behind the pay-cap frontier (Figure 5) and the selection/noise decomposition of the misperception cost. Both are computed firm-year by firm-year from the estimated model, with no

⁶⁶Conservative: skill benchmarks are estimated per spell, so estimation noise biases against cross-spell persistence.

Table 12: Distinguishing adverse selection from competing explanations

Testable margin	Adverse selection predicts	Alternative predicts	Finding
Shape of vote dissent	Two-sided around the model's zero	<i>Power</i> : one-sided, against overpay only	Two-sided; per-unit slopes indistinguishable (Table 4, col. (3))
Location of approval peak	At the model's zero	<i>Peer benchmarking</i> : at the typical contract	$\hat{c} = +0.10$ SD, CI $[-0.31, +0.24]$: contains zero, excludes the typical contract (+0.34); industry-norm distance sign-flips when raced (Fig. 4, App. D.1.2)
Returns of over-perceived firms	No penalty	<i>Power</i> : underperformance (canary)	No penalty in any construction (Table 5)
Relation to entrenchment proxies	Nearly orthogonal	<i>Power</i> : strongly explained by them	Six proxies jointly explain ≤ 2.2 pp of misperception variance; $ \tilde{\epsilon} $ coefficient stable (Table 11)
Behavior at CEO turnover	Persists (base property)	<i>Power</i> : dies with the CEO	Correlation 0.32 across successive CEO spells within firm, against a cross-firm placebo of 0.00 ($z \approx 20$; replication Table T34) ⁶⁶
Incentive slopes around 2018	No break	<i>Tax/162(m) repeal</i> : discontinuity	No break: median slope moves +5% across the boundary with no jump in any year (replication package)

Notes: Each row states the margin on which the adverse-selection account and a named alternative make divergent predictions, and the evidence. No single row is dispositive; jointly, every margin resolves in the adverse-selection direction.

re-estimation, and both reuse the exact accounting of the cost measure in Section 5.5; their consistency is verified numerically in the replication package (the counterfactual code reproduces the Section 5.5 headline cost to three decimal places when evaluated at actual beliefs and contracts).

Inputs. For each firm-year the estimation delivers the true-skill estimate $\hat{b}_i^*$, the belief $\tilde{b}_{it}$ implied by the observed contract, the scale factor ψ_i (the firm's median in-sample market capitalization), the observed contract slope θ_{it} , and lagged market capitalization V_{it-1} ; globally, the estimates $\hat{\sigma}_{FIRM}$ and $\hat{\mu}_{INV}$ and the calibrated risk-aversion parameter $\eta = 4 \times 10^{-4}$. Write $\varsigma_i \equiv \hat{\sigma}_{FIRM} \psi_i$ for firm-scaled fundamental volatility. The CEO's effort-cost parameter is normalized to one (absorbed into the units of skill), as in the estimation.

Value accounting. Under the contracting model, a shareholder base holding belief x about CEO skill chooses slope

$$\theta(x) = \frac{x^2}{x^2 + \eta \varsigma_i^2},$$

and grants base pay according to the participation-constraint formula

$$\phi(s, x) = \omega - s V_{it-1} - \frac{1}{2} s^2 x^2 + \frac{\eta}{2} s^2 \varsigma_i^2,$$

with the level normalization $\omega = 0$. Expected firm value at *true* skill under an arbitrary slope s is $\mathbb{E}[V] = V_{it-1} + s (\hat{b}_i^*)^2$ (effort $a = s \hat{b}_i^*$ evaluated at true skill), and expected compensation is $\phi(s, x) + s \mathbb{E}[V]$. Net shareholder value from a

contract with slope s written under belief x is therefore

$$\text{val}(s, x) = (1 - s) \mathbb{E}[V] - \phi(s, x),$$

and the cost of contracting under belief x rather than the truth is

$$C(x) = \text{val}(\theta(\hat{b}_i^*), \hat{b}_i^*) - \text{val}(\theta(x), x) \geq 0.$$

Evaluated at $x = \tilde{b}_{it}$ with the observed slope, C is identical to the cost measure of Section 5.5. All dollar quantities are converted to 2024 dollars using the January CPI-U and winsorized at the 1st and 99th percentiles, matching the treatment of the cost tables.

The cap experiment. A uniform cap $\bar{\theta}$ replaces each observed slope with $\theta_{it}^{cap} = \min(\theta_{it}, \bar{\theta})$, leaving beliefs unchanged, so base pay re-adjusts to $\phi(\theta_{it}^{cap}, \tilde{b}_{it})$. The per-observation effect is

$$\Delta_{it}(\bar{\theta}) = \text{val}(\theta_{it}^{cap}, \tilde{b}_{it}) - \text{val}(\theta_{it}, \tilde{b}_{it}),$$

which is identically zero for unconstrained firm-years. The frontier evaluates $\bar{\theta}$ on a sixty-point grid of percentiles of the observed slope distribution (0.5 through 99.9, with finer spacing in the upper tail), aggregating Δ_{it} across all firm-years and dividing by the 22 sample years to express effects in 2024 dollars per year. The “gains” and “losses” series in Figure 5 sum the positive and negative Δ_{it} separately; a bound firm-year gains when the cap moves its slope toward $\theta(\hat{b}_i^*)$ by less than the initial distortion, and loses when the cap pushes the slope below it.

The decomposition. The cost of misperception is decomposed by re-contracting under counterfactual beliefs: $x = \hat{b}_i^* + \hat{\mu}_{INV}$ isolates the systematic adverse-selection component, and $x = \hat{b}_i^* + (\tilde{b}_{it} - \hat{b}_i^* - \hat{\mu}_{INV})$ isolates the idiosyncratic component. In each case the slope is re-derived as $\theta(x)$ and the cost $C(x)$ recomputed. The two components do not sum exactly to the total because C is convex in the belief error—noise centered at $\hat{\mu}_{INV}$ is more expensive than noise centered at zero—and the shortfall is reported as an interaction term.

Uncertainty. Confidence bands re-run the entire exercise—grid, aggregation, decomposition—inside each of the 1,000 cluster-bootstrap replications of the structural estimation, using that replication’s firm-level estimates $(\hat{b}_i^*, \tilde{b}_{it})$ and global parameters $(\hat{\sigma}_{FIRM}, \hat{\mu}_{INV})$. Reported intervals are pointwise 2.5th/97.5th percentiles across replications; the interval for the optimal cap’s location is the percentile interval of the per-replication arg max.

Scope and caveats. Two limitations are inherent to the design. First, the experiment is defined on the incentive slope, not on dollar pay levels: the level of base pay is a free normalization in the model (ω is not identified), so the estimates

can price a ceiling on variable-pay intensity but not a dollar cap on total compensation. Second, the exercise is one of partial equilibrium: beliefs, the composition of the shareholder base, and CEO participation are held fixed under the policy, and effort responds only through the contract ($a = s \hat{b}_i^*$). A cap that changed who selects into the stock, or that moved CEOs across firms, would generate effects outside the model. The replication package implements the computation in a single script (`make_counterfactual_outputs.py`), including the internal consistency check against the Section 5.5 cost measure.

The floor experiment. Figure 6 reports the mirror-image experiment: a uniform *floor* $\bar{\theta}$ under which each firm-year's slope becomes $\max(\theta_{it}, \bar{\theta})$, with all other machinery identical to the cap. The frontier is one-sided in the opposite sense: gross gains never exceed roughly \$0.04 billion per year at any floor level, so the net effect is dominated by losses almost immediately, and the optimum is economically negligible. As discussed in Section 7, the asymmetry between the two instruments is informational—the upper tail of the slope distribution is a usable signal of over-perception, while the lower tail predominantly reflects genuinely low optimal slopes—so the floor's failure is an instance of, rather than an exception to, the paper's observability boundary.




Figure 6: Estimated welfare effects of a uniform floor on incentive pay: the mirror image of the cap experiment of Figure 5, replacing each firm-year's slope with $\max(\theta_{it}, \bar{\theta})$. The optimal floor recovers essentially nothing (about \$5 million per year; bootstrap interval [\$1M, \$8M]), floors above roughly the eighth percentile destroy value on net, and gross gains never exceed about \$0.04 billion per year at any floor level; the axes are truncated—the net effect declines steeply beyond the plotted range, reaching $-\$17$ billion per year by the 40th percentile. The failure is informational: a low contract slope predominantly marks firms whose optimal slope is genuinely low, while under-perceived firms sit in the middle of the slope distribution, so no uniform threshold in contract space can locate them. Construction, units, and bootstrap procedure as in Figure 5.

D.5 Mechanism-Discipline Tests: CEO Retention and Contract Self-Correction

This appendix reports two auxiliary panel tests. Both are run on the estimation sample at the firm-year level, keyed to the same annual quintiles of standardized misperception, $\hat{\epsilon}_{it}^{INV} / (\hat{\sigma}_{FIRM} \psi_i)$, used by the asset-pricing sorts, with the aligned middle quintile (Q3) as the omitted category. Both control for $\ln(\text{market capitalization})$ and year fixed effects, with standard errors clustered by firm.

CEO retention. The first column of Table 13 estimates a linear probability model of CEO turnover: the dependent variable equals one if the firm’s CEO spell identifier changes between t and $t+1$, conditional on the firm appearing in the estimation sample in $t+1$ (the last observed year of each firm is necessarily censored). The base rate is 7.5 percent. If the directional return pattern operated through a retention channel—under-appreciated CEOs departing, with underperformance reflecting succession disruption—turnover should be elevated in the under-perceived quintiles. It is not: the point estimate for the most under-perceived quintile is *negative* (2.4 percentage points below the aligned middle, $t = -1.7$). The one robust turnover result runs on the other side: CEOs of the most over-perceived firms are 4.2 percentage points *less* likely to depart ($t = -2.4$), an entrenchment pattern consistent with the absence of belief-driven dissent documented in the say-on-pay analysis.

Contract self-correction. The second column examines whether the contract distortion persists, using same-CEO growth in log total compensation from t to $t+1$ (winsorized at the 1st and 99th percentiles; the same-CEO restriction ensures the change reflects re-contracting rather than succession). Pay growth is steeply and monotonically decreasing in current misperception: CEOs in the most under-perceived quintile receive pay growth 35 log points above the aligned middle ($t = 23$), and those in the most over-perceived quintile 44 log points below ($t = -21$). Part of this reversion is mechanical—transitory components of measured misperception revert by construction, and compensation reverts with them—so the estimates should be read as documenting the *speed* of reversion rather than a causal response. The relevant implication for the mechanism discussion is timing: the contract misalignment associated with a given firm-year is largely corrected within one year, so the directional pattern of Appendix E.1.1 accrues *during* the correction window, not under a persistent distortion. An explanation for the penalty must therefore operate through the one-year window in which the flat contract was in force—as the incentive-channel account does—rather than through cumulative long-run misalignment.

These tests are deliberately modest in ambition: they discipline candidate mechanisms rather than identify one. Their design, code, and the underlying quintile assignments are part of the replication package (`make_mechanism_outputs.py`), and the reported coefficients are verified by the package’s consistency checks.

Table 13: Mechanism-discipline tests: CEO retention and contract self-correction

	(1) CEO turnover _{t+1}	(2) $\Delta \ln(\text{pay})_{t+1}$, same CEO
Q1 (most under-perceived)	-0.0236* (-1.71)	+0.3466*** (+23.16)
Q2	-0.0026 (-0.20)	+0.1542*** (+14.86)
Q4	-0.0201 (-1.59)	-0.1408*** (-13.95)
Q5 (most over-perceived)	-0.0424** (-2.37)	-0.4404*** (-21.13)
Observations	27,190	27,083
Mean dep. var.	0.075	0.092
ln(mktcap), year FE	Yes	Yes

Notes: Firm-year panel regressions on the annual quintiles of standardized misperception $\hat{\epsilon}_{it}^{INV} / (\hat{\sigma}_{FIRM} \psi_i)$ used by the asset-pricing sorts ($N\mu$ specification; aligned middle quintile Q3 omitted). Column (1): linear probability model of a CEO-spell change between t and $t+1$, conditional on the firm appearing in the estimation sample in $t+1$ (base rate 0.075). Column (2): same-CEO growth in log total compensation (tdc1) from t to $t+1$, winsorized at the 1st/99th percentiles. Both columns include ln(market capitalization) and year fixed effects; t -statistics from standard errors clustered by firm (PERMCO) in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

E Appendix: Robustness Checks

E.1 Alternative Regression Specifications

The reduced-form association between active ownership and incentive-heavy compensation (Table 1) is not specific to CEO pay. Table 14 re-estimates all five specifications on the full executive-year panel, measuring variable compensation for every top-5 named executive rather than the CEO alone ($N \approx 84,000$). The results are close to the CEO-sample baselines: contemporaneous active ownership carries a coefficient of 0.127 ($t = 6.1$; CEO sample: 0.123), and lagged ownership of 0.108 ($t = 5.1$). One column reads differently and deserves the caveat: in the leads-and-lags specification, the one-year *lead* of active ownership is marginally significant in this much larger panel (0.057, $t = 2.4$), reflecting the persistence of ownership shares splitting across adjacent-year terms. The reverse-causality inference in the main text therefore rests on the CEO-sample specification of Table 1, where the lead is cleanly insignificant; the appendix table's role is robustness of the *level* association across the executive team, which it delivers.

E.1.1 The Directional Return Pattern

The quintile alphas of Table 5 rise from the under-perceived to the over-perceived end of the sort (from -0.17% to $+0.23\%$ per month around a sample-universe alpha of roughly $+0.15\%$). This appendix documents the directional Q1–Q5 spread that summarizes the pattern, its robustness battery, and why the paper reports it as an open question. The spread is self-financing, so alphas are estimated on raw spread returns:⁶⁷

$$r_t^{LS} = \alpha + \beta' F_t + \varepsilon_t,$$

⁶⁷ Subtracting the risk-free rate from a zero-investment position, as in earlier drafts, overstates the magnitude of alpha by roughly the mean risk-free rate, about 0.1% per month.

Table 14: Reduced Form Robustness: All Top-5 Executives

	<i>Dependent variable: Variable Compensation as a Percent of Total Compensation (all top-5 executives)</i>				
	Baseline (1)	Baseline Lag (2)	Arbs (3)	CEO Tenure (4)	Leads and Lags (5)
Constant	0.182*** (0.025)	0.223*** (0.028)	0.206*** (0.026)	0.270*** (0.039)	0.234*** (0.032)
Active Investor Percent (act_io)	0.127*** (0.021)		0.098*** (0.021)		
ln(Market Cap)	0.047*** (0.003)	0.046*** (0.003)	0.047*** (0.003)	0.044*** (0.005)	0.042*** (0.003)
Bid-Ask Spread			0.000 (0.000)		
Yearly Avg. Trading Vol.			0.000 (0.000)		
Stock Price Volatility			0.000*** (0.000)		
CEO Tenure (yrs.)				0.000 (0.000)	
act_io _{t+1}					0.057** (0.024)
act_io _{t-1}		0.108*** (0.021)		0.130*** (0.029)	0.074*** (0.023)
Year FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Observations	84,027	73,732	81,705	20,555	63,477
R ²	0.530	0.542	0.530	0.599	0.560
Adjusted R ²	0.518	0.529	0.518	0.557	0.547

Notes: OLS estimates on the executive-year panel (all top-5 named executives), 2013–2022 (ExecuComp × FactSet × CRSP). Dependent variable: variable compensation as a share of total compensation, (tdc1 – total_curr)/tdc1. Active Investor Percent is total institutional ownership minus top-3 passive ownership; continuous regressors winsorized at the 1st/99th percentiles. Lagged and lead ownership are one-year firm-level lags/leads of act_io. Standard errors clustered by firm (PERMCO) in parentheses. *p<0.1; **p<0.05; ***p<0.01.

with F_t the Fama–French three- or five-factor set. Table 15 reports the results: -0.42% per month under FF3 ($t = -2.04$) and -0.40% under FF5 ($t = -1.79$), with cluster-bootstrap intervals of $[-0.89\%, -0.08\%]$ and $[-0.88\%, -0.08\%]$ that exclude zero.

Table 15: Fama–French regressions of the misperception long–short portfolio

	FF3	FF5	Small FF5	Med. FF5	Large FF5	V FF5	Orth. FF5
α (monthly)	-0.0042** (-2.04) [-0.0089, -0.0008]	-0.0040* (-1.79) [-0.0088, -0.0008]	0.0023 (1.02)	-0.0067*** (-4.16)	-0.0062*** (-3.29)	0.0019 (1.27)	-0.0072*** (-3.04)
MKT	0.076 (1.61)	0.053 (1.06)	-0.176** (-2.22)	0.023 (0.48)	0.128** (2.33)	-0.127*** (-3.75)	0.132** (2.00)
SMB	-0.556*** (-6.39)	-0.567*** (-5.83)	-0.205** (-2.07)	-0.035 (-0.48)	-0.158* (-1.71)	-0.221*** (-3.20)	-0.132 (-1.11)
HML	-0.480*** (-6.28)	-0.479*** (-6.08)	-0.110 (-1.26)	-0.232*** (-3.60)	-0.402*** (-4.95)	0.154* (1.94)	-0.435*** (-5.27)
UMD		-0.147*** (-3.12)	0.044 (0.80)	-0.106*** (-2.73)	-0.142*** (-2.75)	0.063 (1.03)	-0.131** (-2.00)
RMW		0.050 (0.34)	-0.012 (-0.09)	-0.027 (-0.24)	0.051 (0.38)	-0.028 (-0.41)	0.051 (0.29)
CMA		0.120 (0.80)	-0.046 (-0.28)	-0.061 (-0.66)	0.050 (0.37)	0.158 (1.32)	0.228 (1.47)
Months	258	258	258	258	258	258	258
R^2	0.365	0.403	0.140	0.152	0.217	0.267	0.149

Notes: Long–short portfolio: long the lowest and short the highest quintile of $\hat{\varepsilon}_{t-1}^{INV} / (\hat{\sigma}_{FIRM} \psi_t)$ (annual sort, value-weighted, July–June holding). $N\mu$ specification. HAC(12) t -statistics in parentheses. Square brackets under α : 95% cluster-bootstrap confidence intervals from 1,000 structural re-estimations. Size terciles formed annually on market capitalization. The V column is the symmetric-misperception portfolio (long the aligned middle quintile Q3, short one-half each of Q1 and Q5); its null alpha shows misalignment *per se* is unpriced after factor adjustment. The Orth. column sorts on the component of misperception orthogonal to the observable contract slope (per-year cross-sectional residualization on $\hat{\theta}$), isolating the belief signal that observers of the contract cannot see.

It is not a small-firm artifact: within size groups the spread is -0.67% ($t = -4.16$) among mid-caps and -0.62% ($t = -3.29$) among large caps, while absent among small caps, and the size-conditional construction delivers -0.50% per month ($t = -2.97$; bootstrap interval $[-0.81\%, -0.21\%]$). It is stable across subperiods (size-conditional spread of -0.62% , $t = -2.89$ before 2013 and -0.43% , $t = -2.23$ after) and across distributional specifications (the truncated-normal estimates deliver -0.40% , $t = -3.14$, in the same construction). Sorting on the component of misperception orthogonal to the observable contract slope yields a stronger spread of -0.72% per month ($t = -3.04$; Table 15, Orth. column), so the pattern is not a repackaging of observable contract steepness.

Two considerations nonetheless counsel restraint. First, no individual quintile on the under-perceived side survives a Bonferroni correction, so the pattern is a property of the joint distribution rather than of an identifiable leg. Second, the real-time expanding-window constructions, while uniformly negative in point estimate, are statistically indistinguishable from zero (Table 16: spread alphas of -0.26% to -0.27% per month, $|t|$ between 1.1 and 1.6). The attenuation is informative but not disqualifying: the expanding-window confidence interval (approximately $[-0.64\%, +0.12\%]$) contains the full-sample point estimate, and attenuation of exactly this form is what errors-in-variables predicts when early estimation vintages sort firms with noise. Moreover, the look-ahead embedded in the full-sample sort runs *against* the observed sign: a firm about to perform well receives a higher skill estimate and is sorted toward the under-perceived

quintile, a mechanical bias toward *positive* alpha in that leg—the opposite of what is found. What the real-time attenuation does establish is that the pattern was not exploitable as a trading strategy with vintage estimates, a claim the paper does not make.

Table 16: Expanding-window and double-sorted Fama–French regressions

	FF3	FF5	Double-sorted FF5	V FF5	V Double-sorted FF5
α (monthly)	-0.0027 (-1.55)	-0.0026 (-1.34)	-0.0021 (-1.08)	0.0003 (0.20)	-0.0015 (-0.99)
MKT	0.040 (0.90)	0.016 (0.35)	0.169*** (3.10)	-0.148*** (-3.83)	-0.140*** (-4.90)
SMB	-0.609*** (-6.40)	-0.651*** (-6.25)	-0.244*** (-2.97)	-0.053 (-1.01)	-0.084 (-1.52)
HML	-0.616*** (-8.26)	-0.531*** (-4.97)	-0.486*** (-7.57)	0.052 (0.94)	0.024 (0.38)
UMD		-0.113** (-2.12)	-0.038 (-0.85)	0.004 (0.06)	0.113*** (2.71)
RMW		0.001 (0.00)	-0.104 (-0.73)	0.115 (1.45)	0.038 (0.52)
CMA		-0.013 (-0.07)	-0.004 (-0.03)	0.266*** (2.66)	0.269*** (2.88)
Months	198	198	198	198	198
R^2	0.539	0.562	0.395	0.268	0.403

Notes: Expanding-window construction: ε^{INV} re-estimated each year using only data through year t ; portfolios formed in July of $t+1$. Double sort: size quintiles first, misperception quintiles within. $N\mu$ specification, HAC(12) t -statistics in parentheses. V columns: the symmetric-misperception portfolio (long aligned middle, short both tails equally) under the same real-time constructions.

Why might any pattern persist? A natural objection is that a sophisticated investor could invert the observed contract slope to recover the shareholder base’s belief—the inversion this paper’s estimation performs. Two answers. First, the inversion requires the structural apparatus: a contracting model, an estimate of fundamental volatility, and a performance-disciplined benchmark for true skill; the contract does not wear its belief content on its face. Second, Section 6.1 shows that no observed market participant performs it: the belief component of contracts moves no votes for any electorate. As for mechanism, the candidate explanations—belief-correlated information incorporating slowly into prices in the tradition of Miller (1977) and Hong and Stein (2007), or incentive effects of contract misalignment—are observationally equivalent in these data. The test that would separate them, whether the spread concentrates at information events such as earnings announcements, requires daily return data and is left for future work. I flag the pattern rather than build on it.

E.2 Identification of the Structural Model

Sources of identifying variation. Three moments jointly identify the parameters.

1. **Contract mapping moment:** Observed θ_{it} and a candidate σ_{FIRM} yield the implied $\tilde{b}_{it}(\sigma_{FIRM})$, providing a functional relationship between σ_{FIRM} and the cross-sectional dispersion of perceived skill.
2. **Belief distribution moment:** Repeated observations of $\tilde{b}_{it}$ within a firm identify the dispersion σ_{INV} and drift

μ_{INV} of belief errors around the firm-specific skill b_i^* . The *location* b_i^* is not identified by this moment alone: a level shift in candidate skill is absorbed one-for-one by the estimated errors.

3. **Performance moment:** The residuals ϵ_{it}^{FIRM} from realized performance identify σ_{FIRM} as the variance of fundamental shocks consistent with the observed distribution of $\Delta V_{it} - \theta_{it}(b_i^*)^2 - c_{it}(V_{it})$ —and, critically, anchor the location b_i^* : candidate skill levels that rationalize the contract history but are inconsistent with realized performance are penalized. It is this moment that pins the zero point of misperception and thereby makes μ_{INV} a meaningful estimate rather than a normalization.

Cross-firm heterogeneity strengthens identification because σ_{INV} , σ_{FIRM} , and μ_{INV} must jointly rationalize variation in both contracts and performance across a large number of firms with distinct b_i^* .

Potential threats to identification.

1. **Scale indeterminacy:** If θ_{it} and ΔV_{it} are measured in inconsistent units, σ_{FIRM} can be rescaled arbitrarily. Ensuring that $\theta_{it}(b_i^*)^2$ and ΔV_{it} are expressed in the same monetary units is crucial.
2. **Weak variation:** Limited variation in θ_{it} or highly noisy ΔV_{it} flattens the likelihood surface, leading to weak identification of σ_{FIRM} and σ_{INV} .
3. **Confounding between σ_{FIRM} and b_i^* :** Because σ_{FIRM} enters both the $\theta \rightarrow \tilde{b}$ mapping and the performance variance, misspecification (e.g., time-varying σ_{FIRM} or measurement error in θ) can bias estimates.
4. **Distributional misspecification:** The parametric form of the belief-error distribution encodes a specific mechanism for investor misperception. If selection operates differently (e.g., driven by performance shocks rather than perceived skill), the functional form is misspecified and identification may fail. Section E.3 assesses sensitivity to this choice directly.

Three exercises reported in the remainder of this appendix support identification in practice: the Monte Carlo experiments of Section E.6 show that the estimator recovers all three parameters at the production sample size; the multistart analysis in the same section shows the likelihood has a unique, well-conditioned optimum on the actual panel; and the alternative-distribution estimates of Section E.3 show which conclusions do and do not depend on the parametric form.

E.3 Alternative Error Distributions

The cost estimates are robust to the assumed distribution of investor errors. Table 17 reports point-estimate costs under three specifications computed with identical methodology: the truncated-normal specification of earlier drafts (one-sided errors, overvaluation imposed by the support), a log-normal alternative, and the normal-with-drift ($N\mu$)

specification adopted in the main text. The sign of the cost is invariant across all three; the positive-support specifications produce magnitudes roughly twice the $N\mu$ values because they pin the firm-level skill estimate at its boundary. Table 18 reports the truncated-normal parameter estimates with bootstrap confidence intervals for comparison with the main specification.

Table 17: Cost estimates under alternative ε^{INV} distributions

Specification	Mean (2024\$M)	Median (2024\$M)	Aggregate (2024\$B)
Truncated normal	-2.32	-0.146	-67.2
Log-normal	-1.98	-0.098	-57.4
Normal (μ_I free)	-1.16	-0.067	-33.6

Notes: Point-estimate annual costs under identical fixed- φ methodology. The log-normal likelihood is maximized at a near-degenerate configuration that reproduces the truncated-normal boundary solution (see text); its role is to confirm that the cost is identified by the firm-side moment condition rather than the assumed shape of investor errors.

Table 18: Structural parameter estimates (truncated-normal robustness)

Parameter	Estimate	95% CI
σ_I	2,944	[1,634, 4,629]
σ_F	0.538	[0.364, 0.769]

Notes: Maximum-likelihood estimates for the TN specification with 95% confidence intervals from 1,000 cluster bootstrap replications (firms resampled with replacement, full re-estimation per replication). Pairs-bootstrap intervals (firm-years resampled) are similar throughout, agreeing within 5–15%.

The two alternatives play different roles. The truncated-normal specification is a genuine competing model—it is the specification under which earlier drafts of this paper were written—and its estimates are directly comparable to the main results. The log-normal specification instead functions as a confirming degenerate case: its maximum-likelihood estimate drives the log-scale parameter toward zero, collapsing the belief distribution toward a point mass. That degeneracy is informative rather than embarrassing: it indicates that the *shape* of the belief-error distribution is only weakly identified by the belief moment, and that the cost estimates are disciplined primarily by the firm-side performance moment—which is exactly why the cost magnitudes in Table 17 agree across such different distributional assumptions.

A separate battery perturbs the firm side of the model and the construction of the performance measure one dimension at a time, re-estimating the full model under the current pipeline in each case; the baseline is $(\hat{\sigma}_{INV}, \hat{\sigma}_{FIRM}, \hat{\mu}_{INV}) = (1986, 0.538, 85.2)$.

Positive firm-side drift: centering the performance shock at the sample-average growth rate, $\epsilon_{it}^{FIRM} \sim \mathcal{N}(0.0717\psi_i, \psi_i^2\sigma_{FIRM}^2)$ (also $0.099\psi_i$), leaves the dispersions essentially unchanged ($\hat{\sigma}_{INV}$ 1977–1986, $\hat{\sigma}_{FIRM}$ 0.536–0.538) while $\hat{\mu}_{INV}$ rises to 112 and 119 respectively: drift and optimism trade off along the performance moment exactly as the identification discussion predicts, and the adverse-selection finding is, if anything, strengthened when secular growth is conceded to fundamentals.

Performance measured by total shareholder return: replacing the market-capitalization change ΔV with the firm’s annual total shareholder return times lagged capitalization—which strips net issuance and buybacks out of the perfor-

mance measure—yields (1885, 0.511, 80.3), each parameter inside the baseline’s 95% cluster-bootstrap interval.

Share-class aggregation: consolidating multi-class firms across share classes yields (2188, 0.545, 84.7). *Fiscal-calendar alignment:* restricting to the roughly three-quarters of firm-years with December fiscal year-ends yields (1270, 0.400, 54.9), with $\hat{\mu}_{INV}$ positive and all parameters within the full-sample intervals. (Tables T17–T19 and T21 of the replication package.)

Two profile-likelihood exercises push the drift question further by treating the maintained zero-drift assumptions as points on a grid rather than fixtures. Profiling the firm-side drift over $m \in [-0.05, 0.099]$ (in units of ψ_i) yields a likelihood that is monotone toward *negative* drift, with no interior optimum in the economically sensible range: freed, the parameter chases the implication that the average firm sheds five percent or more of its value annually absent CEO skill, contradicting the sample’s average growth of roughly seven percent. Profiling instead a common drift in CEO skill, $b_{it}^* = b_i^* (1 + \gamma)^{\text{tenure}_{it}}$ over $\gamma \in [-0.05, 0.05]$ per year of tenure, yields a likelihood *minimized* at the constant-skill baseline, improving in both directions. Neither preferred drift is interpretable as its label: within a CEO spell tenure coincides with calendar time, so both parameters largely proxy for secular time-structure that the model does not otherwise carry, and I therefore fix both at their theory-imposed values and read the profiles as bounding sensitivity. The bound is one-sided in the direction that matters: $\hat{\mu}_{INV}$ is positive at every point of both grids—ranging from 48 to 119 along the firm-drift profile and rising to between 173 and 425 along the skill-drift profile—and the baseline specification is the most conservative point of the entire battery, with every alternative raising the estimated optimism. (Tables T23–T24 of the replication package.)

E.4 Sensitivity to the Risk-Aversion Calibration

The estimation calibrates CARA risk aversion at $\eta = 4 \times 10^{-4}$ (per thousand dollars), following Yan et al. (2025); standard calibrations in the contracting literature are stated in relative-risk-aversion terms—typically 2 to 3—which convert to absolute risk aversion via $\eta \approx \text{RRA}/\text{wealth}$ and imply values within the grid examined here at representative CEO wealth.⁶⁸ Table 19 re-estimates the full model at $\eta \in \{1, 2, 4, 8, 16\} \times 10^{-4}$. The results sort into two tiers. Identification-level conclusions are calibration-invariant: $\hat{\sigma}_{INV}$ scales exactly as $\sqrt{\eta}$ (a units effect), $\hat{\sigma}_{FIRM}$ is unchanged to the third decimal, $\hat{\mu}_{INV}$ is positive throughout, and the standardized sorting variable’s ordering is essentially unaffected (Spearman correlations above 0.97), so the portfolio sorts, the say-on-pay results, and the kink-location test do not depend on η . Dollar magnitudes, by contrast, scale approximately linearly in η —as they must, since the contracting loss consists of the risk premium and excess-effort compensation, both proportional to η at a given contract. The headline cost estimates should accordingly be read as calibration-scaled, bounded within roughly \$0.4 to \$2.4 billion per year over the conventional RRA range, with every counterfactual *ranking* (cap versus floor versus belief-dispersion reduction) preserved because all instruments scale together.

⁶⁸The CARA-normal framework and the risk premium $\frac{\eta}{2} \theta^2 \sigma^2$ follow Edmans et al. (2017).

Table 19: Sensitivity to the CARA risk-aversion calibration η

$\eta (\times 10^{-4})$	$\hat{\sigma}_{INV}$	$\hat{\sigma}_{FIRM}$	$\hat{\mu}_{INV}$	$\hat{\mu}/\hat{\sigma}_{INV}$	Cost (\$B/yr)	Rank corr.
1	993	0.538	29.6	0.030	0.37	0.982
2	1,404	0.538	49.9	0.036	0.75	0.994
4 (baseline)	1,986	0.538	85.2	0.043	1.53	1.000
8	2,809	0.538	144.5	0.051	3.10	0.991
16	3,974	0.538	241.5	0.061	6.36	0.966

Notes: Full re-estimation of the structural model at each η , with the misperception cost recomputed at the matched calibration. $\hat{\sigma}_{INV}$ scales as $\sqrt{\eta}$ (each row is $\sqrt{2}$ times the previous) and $\hat{\sigma}_{FIRM}$ is invariant, as the contract inversion implies; $\hat{\mu}_{INV}$ is positive throughout. Rank corr. is the Spearman correlation of the standardized sorting variable with its baseline values: portfolio sorts, the say-on-pay results, and the kink location are unaffected by the calibration. The dollar cost scales approximately linearly in η , as the contracting loss is composed of the risk premium and excess-effort compensation, both proportional to η at a given contract; counterfactual *rankings* are unaffected because all instruments scale together.

E.5 Alternative Measurement Choices: Contract Slope, Risk Aversion, and Value Timing

This appendix probes three measurement choices underlying the structural estimation: the contract-slope measure θ , the assumption of a common risk-aversion coefficient η , and the use of raw (rather than market-adjusted) changes in firm value. Table 20 re-estimates the model under each alternative and traces the consequences through to the two downstream objects that discipline the estimates: the implied dollar losses and the alignment of the recovered belief errors with say-on-pay votes.

Contract slope. Our baseline θ measures the slope of the CEO’s *annual compensation flow* in firm value, which is the object appearing in the model’s contract equation. A natural alternative is the Core–Guay portfolio delta—the sensitivity of the CEO’s accumulated equity *wealth* (shares plus options) to firm value—from the Coles–Daniel–Naveen data Coles et al. (2006). The two measures are genuinely different (median wealth-delta slope roughly seven times steeper; Spearman correlation 0.47), and the comparison is run on the identical matched subsample, on which the baseline measure reproduces the full-sample estimates almost exactly (row 1). The wealth-based slope preserves—indeed amplifies—the paper’s core finding: $\hat{\mu}_{INV}$ remains strictly positive and rises relative to $\hat{\sigma}_{INV}$ (ratio 0.28 versus 0.045). But three coherent diagnostics indicate that the wealth measure is the wrong object for a flow-contract model: the skill estimate is pushed to its zero boundary for 47% of firm–CEO spells (versus 4% at baseline); the implied aggregate loss (\$487B per year) is economically implausible; and—most tellingly—the belief errors recovered under the wealth measure bear *no* relationship to say-on-pay dissent (final column), whereas the flow-based errors predict dissent at $t \approx -5$. Shareholders vote on the annual pay package; that their votes track the flow-based misperception measure and not the wealth-based one is independent, out-of-model evidence for the baseline construction.

Heterogeneous risk aversion. Appendix E.4 shows the level of a common η is not load-bearing. Here I relax commonality: each CEO-year’s CARA coefficient is set inversely proportional to firm-related wealth Coles et al. (2006), median-normalized to the baseline calibration and winsorized at p1/p99—a large perturbation (p90/p10 ratio ≈ 59). The firm-side variance $\hat{\sigma}_{FIRM}$ is unchanged (0.468 in both runs), the standardized bias moves only from 0.046 to 0.051, the sorting variable is largely preserved (rank correlation 0.71), and the recovered errors still predict two-sided dissent ($t = -4.3$). The dollar loss roughly doubles, to \$1.96B per year on the matched sample: because

losses are approximately linear in η and wealth-implied η_i is right-skewed, the homogeneous baseline is, if anything, conservative.

Market-adjusted value changes. Raw ΔV embeds the market’s secular rise, which Appendix E.3’s firm-drift profile will absorb into the drift parameter. Re-estimating with $\Delta V_{adj} = V_t - V_{t-1}(1 + r_{m,t})$ confirms this reading: the log-likelihood profile over the firm-drift parameter, which rises monotonically toward implausibly large positive drift under raw ΔV , is monotonically *decreasing* in drift under the market-adjusted construction—the apparent drift preference is the market return in disguise, not a feature of firm-level fundamentals. All downstream objects are essentially unchanged (rank correlation 0.99; dissent $t = -5.7$; loss \$1.24B versus \$1.53B per year at baseline).

The scaling holds under resampling, not only at the point estimates: a 200-replication cluster bootstrap at $\eta = 8 \times 10^{-4}$ yields $\hat{\mu}_{INV} > 0$ in every replication, a $\hat{\sigma}_{FIRM}$ distribution indistinguishable from the baseline- η bootstrap (medians 0.537 versus 0.527; 95% intervals [0.37, 0.77] versus [0.36, 0.77]), and a $\hat{\sigma}_{INV}$ distribution that is the baseline distribution scaled by $\sqrt{2}$ throughout (scaled medians within 2%; scaled interval [1,078, 3,200] versus [991, 3,220]).

Table 20: Alternative measurement choices: parameters, losses, and vote alignment

Variant	$\hat{\sigma}_{INV}$	$\hat{\mu}_{INV}$	$\hat{\mu}/\hat{\sigma}_{INV}$	Loss (\$B/yr)	$ \tilde{\epsilon} $ dissent (t)	Rank corr.
Baseline θ , CDN-matched sample	1,653	74	0.045	0.81	-0.0172*** (-4.82)	0.30
CDN wealth-delta θ	6,158	1,706	0.277	487.25	-0.0001 (-0.05)	0.30
Wealth-scaled η_i (CDN)	3,456	175	0.051	1.96	-0.0089*** (-4.34)	0.70
Market-adjusted ΔV	1,772	100	0.056	1.24	-0.0188*** (-5.74)	0.99

Notes: Each row re-estimates the structural model under one alternative measurement choice, then recomputes the misperception loss (2024\$, T4/T8 recipe, matched calibration) and the say-on-pay distance regression (Table 4, column (4) specification: approval on $|\tilde{\epsilon}|$ with controls, clustered by firm) using that variant’s recovered belief errors. *CDN wealth-delta θ* : contract slope from the Coles–Daniel–Naveen Coles et al. (2006) Core–Guay CEO portfolio delta (shares plus options) in place of the annual-flow slope; estimated on the identical matched subsample as row 1. *Wealth-scaled η_i* : CARA coefficient inversely proportional to the CEO’s firm-related wealth, median-normalized to the baseline 4×10^{-4} and winsorized at p1/p99 (p90/p10 ratio ≈ 59). *Market-adjusted ΔV* : firm value change net of the calendar-year market return, $\Delta V_{adj} = V_t - V_{t-1}(1 + r_{m,t})$. Rank corr. is the Spearman correlation of the variant’s standardized sorting variable with the baseline. t -statistics clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E.6 Monte Carlo Validation and Optimizer Diagnostics

This section reports two Monte Carlo exercises and two optimizer diagnostics. The Monte Carlo experiments characterize finite-sample behavior of the estimator: a *matched-specification* experiment at $N = 100$ firms, designed to reproduce and diagnose the σ_{INV} bias visible in earlier drafts, and a *scaling* experiment at $N \in \{100, 1,000, 3,693\}$ firms under the main $N\mu$ specification, with 3,693 matching the production panel. In each replication, data are simulated at the estimated parameters and the full estimation pipeline—inner loop over firm-specific skill, outer Nelder–Mead over the global parameters—is re-run from perturbed initializations ($\pm 50\%$ multiplicative perturbation of the starting values).

Table 21 shows that small samples are genuinely hostile to both specifications, through different channels. Under the truncated normal, the median σ_{INV} estimate at $N = 100$ is biased downward by approximately 26%—reproducing, under a matched optimizer, the bias of roughly -23% visible in earlier drafts, and establishing that it is a finite-sample

Table 21: Matched-specification Monte Carlo at $N = 100$ firms

Specification	Parameter	True	MC median	Bias
TruncNormal	σ_{INV}	2,944.0	2,181.5	-25.9%
TruncNormal	σ_{FIRM}	0.538	0.441	-18.1%
Nmu	σ_{INV}	1,986.0	1,459.4	-26.5%
Nmu	σ_{FIRM}	0.538	0.537	-0.2%
Nmu	μ_{INV}	85.2	1,119.5	+1214.0%

Notes: Median estimates across 100 replications, matched data-generating process and estimator. The truncated-normal σ_{INV} bias reproduces the value reported in earlier drafts and is a finite-sample feature of the one-sided-support specification, absent under $N\mu$ at scale (see Table 22).

identification feature of the one-sided-support specification (probability mass forced against the zero boundary), not an optimizer artifact. Under $N\mu$, the same small sample instead produces a location–dispersion trade-off: the median μ_{INV} estimate is more than an order of magnitude too large while σ_{INV} is biased downward by a similar 26%; σ_{FIRM} , disciplined by the performance moment, is essentially unbiased even at $N = 100$.

Table 22: Monte Carlo identification of μ_I by sample size

N firms	μ_I median	bias	σ_I median	bias	σ_F median
100	1,119.5	+1214%	1,459	-26.5%	0.537
1,000	89.1	+5%	1,961	-1.3%	0.565
3,693	90.8	+7%	2,065	+4.0%	0.546

Notes: Median estimates across 100 replications with $\pm 50\%$ multiplicative perturbation of initial values; true values $\mu_I = 85.2$, $\sigma_I = 1,986$, $\sigma_F = 0.538$. The $N = 3,693$ row corresponds to the estimation sample size.

Table 22 shows that these are small-sample phenomena. By $N = 1,000$ firms the median estimates are within roughly 5% of truth on all three parameters, and at the production sample size of $N = 3,693$ the median biases are +4.0% (σ_{INV}), +1.6% (σ_{FIRM}), and +6.5% (μ_{INV}). Medians are the appropriate summary: mean-based bias measures are dominated by a small number of replications in which the simplex collapsed on a boundary.

The convergence flag. Nelder–Mead’s formal convergence criterion deserves comment, because at $N = 100$ only about a quarter of replications formally converge. The flag refers to the simplex contraction tolerance, not to whether a replication recovered the truth: replications that fail to formally converge typically park their best vertex at the correct parameter values while the simplex wanders a flat region without contracting (at $N = 1,000$, the median σ_{INV} across all replications, converged or not, is 2,000 against a truth of 1,986). Three diagnostics confirm the flag is a structural feature of the optimization geometry rather than a tunable threshold: tripling the iteration budget (300 to 1,000) leaves the same replications converged and the same estimates; rescaling all three parameters to order one before optimizing—so that a single tolerance is uniformly meaningful across $\sigma_{INV} \approx 2,000$, $\sigma_{FIRM} \approx 0.5$, and $\mu_{INV} \approx 85$ —changes nothing (estimates identical to 6×10^{-3} across independent runs); and sweeping the function-value tolerance across four orders of magnitude rescues no additional replications. The reported medians are accordingly reproducible to the digit across iteration budgets, scalings, and tolerances.

Multistart on the production panel. Finally, identification at the production sample size is established directly. Re-estimating on the actual $N = 3,693$ panel from six widely separated initializations ($\sigma_{INV} \in \{1000, 1500, 1986, 2940, 3500, 5000\}$, $\sigma_{FIRM} \in \{0.3, 0.5, 1.0\}$, $\mu_{INV} \in \{50, 85, 150, 200, 300, 500\}$), every start converges to the same point— $(\sigma_{INV}, \sigma_{FIRM}, \mu_{INV}) = (1986, 0.538, 85.2)$ —with spread at most 0.001 on each parameter and identical log-likelihood (729,960.60). The likelihood is therefore unimodal on the production panel. The estimate reported in earlier drafts ($\sigma_{INV} = 2,940$, obtained with L-BFGS-B) lies in the same basin but stopped roughly 30 log-likelihood units short of the optimum; Nelder–Mead initialized at that point descends to the same (1986, 0.538, 85.2). The earlier numbers were under-converged rather than located at a different mode. The Monte Carlo experiments above are thus best read as a characterization of small-sample behavior; identification at the sample size actually used is established by the multistart exercise.