Shareholder Voice and Executive Compensation*

John W. Barry†

Rice University

December 28, 2024

Click here for the latest version

Abstract

I estimate a model of CEO compensation with non-binding shareholder approval votes (Say-on-Pay). The Board sets compensation and (relative to shareholders) may prefer high total pay; shareholders can fail the vote and punish the Board for high pay. Failed votes are perceived as costly by the Board and shareholders: Say-on-Pay resembles a costly punishment mechanism, raising firm value by 2.4% on average, despite only 7% of votes failing. I analyze a counterfactual binding vote: failure binds pay to prior levels, which may not reflect current information about CEO ability. The failure rate falls, pay levels increase and firm value decreases.

Keywords: shareholder voice, corporate governance, executive compensation, CEO compensation, shareholder voting, say-on-pay, structural estimation, non-binding shareholder voting, binding shareholder voting

*I thank John Graham, Alon Brav, Simon Gervais, Arthur Taburet and Vish Viswanathan for invaluable guidance, and the entire Finance department at Fuqua for comments. Thank you to discussants Austin Starkweather, Pierre Chaigneau, Oliver Levine, Yuri Tserlukevich and Felix Zhiyu Feng, and to James Pinnington, Alan Crane, Kevin Crotty, Nadya Malenko, Andrey Malenko, Alex Edmans, Josh White, Mike Wittry, Wenyu Wang, Yufeng Wu, and many others for comments. The staff at the DCC, particularly Tom Milledge, were exceedingly helpful in implementing the estimation. All errors belong to I.

†Email: john.w.barry@rice.edu
Web: https://johnwbarry.info/

1. Introduction

Shareholders elect the Board of Directors but the Board need not represent their interests (Shleifer and Vishny, 1997). A primary manifestation of this agency problem is in executive compensation: the Board should set compensation to align the interests of management and shareholders, yet directors may have an incentive to favor the CEO (Bebchuk and Fried, 2003).

When shareholders disagree with compensation decisions and exerting control is not viable, shareholders can convey dissent through *voice* (Hirschman, 1970; Cuñat et al., 2016). Say-on-Pay (SOP), a non-binding shareholder approval vote on CEO compensation policy, is the primary channel through which shareholders can voice dissent. While non-binding, SOP is in essence a vote of confidence on the Board's pay decisions and CEO performance. As compensation is the key tool to limit potential agency problems in managerial decision-making, SOP is a potentially important governance mechanism.

Yet the impact of SOP is unclear. Compensation policies receive over 90% support on average and only about 7% of SOP votes in the US fail.[1] As Figure 1 shows, the generally positive outcomes of SOP may be hard to reconcile with both survey evidence on shareholder views about the level of total CEO pay (Edmans et al., 2023), as well as the literature studying CEO influence on the pay-setting process (e.g., Bertrand and Mullainathan, 2001; Coles et al., 2014).

An important consideration is that these approval votes are *ex post*, endogenous outcomes, occurring after compensation has been set and firm performance realized. What determines the impact of SOP is how much the Board internalizes the cost of failure into *ex ante* compensation decision. As Figure 1 suggests, it may be just as important to understand shareholders' apparent hesitancy to fail SOP votes and dissent from the Board on compensation policy.

This paper's goal is to build a structural model to quantify the influence of SOP on compensation policy and explore the mechanism by which this influence occurs. In the model, the Board decides compensation policy and may be biased towards high total pay. Shareholders hold the SOP: failure punishes the Board for overpayment, yet dissent may be costly to shareholders. Estimating the model will assess how much boards internalize the cost of SOP failure into pay decisions and whether shareholders consider SOP failure a costly outcome, quantifying fully the influence of SOP on compensation policy.

[1] SOP was formalized in the US as part of Dodd-Frank in 2010. SOP proposals are put forth by management at the annual shareholder meeting, and shareholders vote on the CEO's compensation from the previous fiscal year. I use "failure" to refer to SOP proposals that do not garner the required support from shareholders. In the US, SOP votes are non-binding, so there is no threshold which forces the Board to change pay policy. However, the understood threshold for SOP failure is 70% support (or 30% voting against, see Dey et al., 2024; ISS, 2022, Section 5 "Compensation"), though 50% is also an important threshold (Hauder, 2019). Further, SOP votes are *ex post* approval votes on the previous year's CEO compensation, not advisory votes on *proposed* compensation.

Figure 1. Say-On-Pay results and shareholder satisfaction with CEO pay

Panel A displays SOP passing rate in the US by year from 2010 to 2020 is about 93% (a passing vote garners over 70% support, ISS, 2022). Panel B displays results from a survey question in Table OA4 of Edmans et al. (2023) that asks UK institutional investors about the levels of the CEOs' pay; over 75% of survey respondents state that they believe CEO pay is too high.

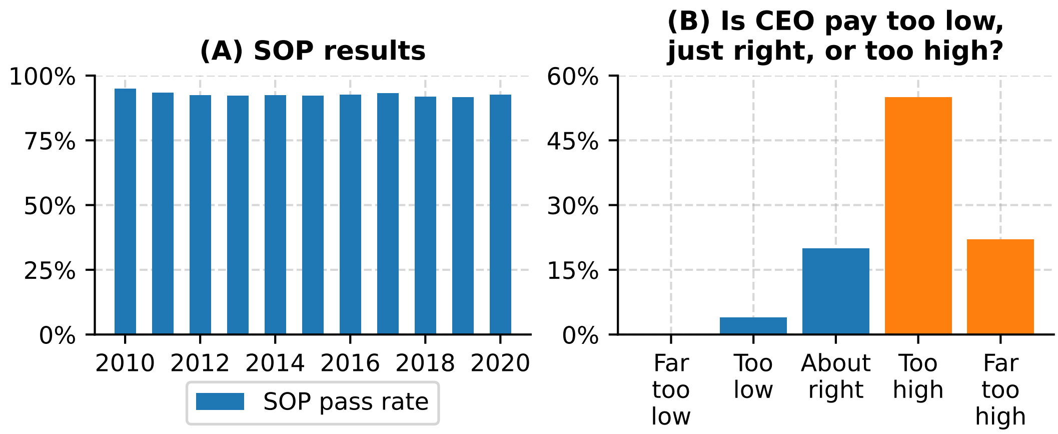

However, factors beyond these potential costs also influence compensation and voting decisions. The size (or even existence) of the Board's bias is not obvious and may vary across the CEO's tenure (e.g., by potentially capturing the Board over time, Coles et al., 2014). Some CEOs are more skilled than others and will receive different compensation for their effort. The Board and shareholders cannot observe CEO skill directly and may have different beliefs about the CEO's ability (Taylor, 2010); public and private information will influence learning, and consequently, pay and vote decisions. Separately quantifying the effects of these forces is necessary to fully understand the impact of SOP.

I estimate model parameters via indirect inference and the model matches key features of the data. The model replicates the observed SOP failure rate very closely. Importantly, it matches the sensitivity of SOP failure likelihood to both the level of total pay and company performance, the primary determinants of SOP vote outcomes (Fisch et al., 2018).

The estimation produces several results. To start, boards shift a share of surplus towards the CEO (relative to the profit-maximizing wage), which I refer to as *board capture*. I estimate that the average S&P1500 CEO receives 16.4% of expected surplus in the first year of tenure, and this grows by about 6.6% per year (reaching 58.1% by year 20 of tenure). This is in line with Taylor (2013), who finds that CEOs capture about half of the surplus from positive updates about their ability.

How does CEO surplus capture of this magnitude square with the seemingly low SOP failure rate? My structural model provides an answer, considering the costs internalized by directors and shareholders from failed SOP votes. First, for SOP to impact compensation policy, it must be that the Board internalizes the threat of vote failure into their decision. To explain observed behavior, I estimate that

Boards internalize a utility cost from SOP failure that is equivalent to 3.14% of firm value.[2] While the unconditional failure rate of about 7% means the cost is closer to 0.21% of value in expectation, the threat of costly SOP failure disciplines CEO compensation, even when failure is *ex ante* unlikely. I estimate that SOP as a disciplining mechanism brings total pay levels down by 6.6% on average, in line with Correa and Lel (2016), who find the adoption of SOP brought wages down by about 7%. Hence, *shareholder voice affects executive compensation policy.*

Second, failed SOP votes are considered costly by shareholders. I estimate that shareholders in the model internalize a cost to SOP failure equivalent to 0.90% of value (about 0.05% in expectation).[3] This aligns with survey evidence from Edmans et al. (2023): shareholders state that failing the SOP may be undesirable, for example because they are hesitant to dissent from the Board on a prominent policy.[4] Though SOP failure is conditionally costly, my estimates suggest that the disciplining effect of SOP improves firm value by 2.35% on average, consistent with Cuñat et al. (2016), who find that voting to adopt SOP as a governance tool increased market value by 5% for firms close to the majority threshold.

These results highlight the simple economics through which these votes impact CEO pay policy and value: SOP resembles a costly punishment mechanism (Silveira, 2017). The threat of punishment disciplines the Board, even in states when SOP failure is unlikely; this discipline brings the Board's CEO compensation policy (and consequently, the CEO's action) closer to that which maximizes shareholder value. Thus, providing shareholders with the regularly-occurring punishment technology is value-enhancing, even if punishment is costly and rarely occurs.

To infer the magnitude of unobservable model parameters, the structural estimation uses observed, endogenous patterns in company performance, CEO pay and SOP vote outcomes. Its success hinges in part on whether there are sensible patterns to reinforce the structural results. As described next, I document several empirical facts about the Board and shareholder failure costs which underpin the model.

The first set of new descriptive facts shows that SOP failure may lead to negative effects for directors, in support of the magnitude of the Board's cost. I find that failing SOP votes is a career and reputation concern for directors: SOP failure is associated with a 2 percentage point (pp) increase in the likelihood that a compensation committee director leaves or is removed from the Board (a 20% larger likelihood of turnover relative to the non-fail group). For directors that remain on the board, I find they are more likely to be off the compensation committee: SOP failure is associated with a 1.5 pp increase in the

[2]Because the Board maximizes a mix of firm value and CEO utility, the mix dictated by board capture, this cost can be equivalently stated as a percentage of the Board's equity stake in the firm

[3]That is, the cost is equivalent to 0.90% (0.05% in expectation) of each shareholder's equity stake in the firm.

[4]See Online Appendix A of Edmans et al. (2023) for discussion of potential costs from shareholder dissent in SOP.

likelihood they are off the compensation committee the next year (26% larger compared to the baseline). Interestingly, I find that failed SOP votes lead to *external* reputational damage for directors. A failed SOP at a director's current firm is associated with a decrease in outside Board positions at other firms (a 2 pp increase in the likelihood that a director loses at least one outside board position). This evidence is in line with Fos and Tsoutsoura (2014) and Aggarwal et al. (2019, 2023); however to the best of the knowledge, my paper is the first to document internal and external costs to directors tied to SOP failure.

While directors generally wish to be re-appointed to the Board, I argue that a large portion of the Board's perceived failure cost acts through a prestige channel (Fos and Tsoutsoura, 2014; Bebchuk and Fried, 2003). SOP failure is a public negative performance evaluation from shareholders on a prominent policy. Directors may have an incentive to favor the CEO; however, the threat of SOP failure pushes their incentives towards shareholders.

My estimation also shows that SOP failures are perceived as costly by shareholders. In surveys (Edmans et al., 2023), shareholders state they avoid failure to maintain relations with the Board. The negative performance evaluation aspect of SOP failure may commit shareholders to raising the rate of Board turnover, which is not costless.

As further evidence consistent with this cost, I show there is bunching in SOP vote outcomes: I uncover excess density directly below the failure thresholds of 30% and 50%.[5] In the data, bunching is consistent with shareholders internalizing a cost from SOP failure. Blockholders (often pivotal in SOP votes) may have an incentive to force a close pass relative to a close fail, precisely because they internalize a cost to SOP failure. Similar bunching evidence was first documented in Listokin (2008), and has been found in Babenko et al. (2019) in the broader context of management proposals, and in Bach and Metzger (2019) for shareholder proposals.

This empirical evidence provides further insight about the economics at play. SOP votes are non-binding, so they do not directly impact the Board's compensation policy. To give power to voice, shareholders hold the Board accountable when the SOP vote fails; for example, by exerting control and replacing directors in the future. Failure is thus costly for the Board, but not a free ride for shareholders.

These findings further our understanding of shareholder voting. I show that the threat of a failed vote may alter corporate decisions *ex ante*, resulting in a passing vote *ex post*. A direct consequence is that high shareholder support in management proposals does not imply ineffective monitoring. Moreover, observed voting decisions may not reveal shareholder preferences — one must consider how off-equilibrium vote outcomes influence corporate policies and consequently, equilibrium voting decisions.

[5]In SOP, 30% and 50% of shareholders voting against the SOP are key thresholds (see Hauder, 2019).

In the last part of the paper, I use the structural model to analyze a policy-relevant counterfactual implementation of Say-on-Pay. While my estimates indicate that non-binding SOP is an effective governance mechanism, it remains unclear whether a *binding* shareholder vote would enhance this efficacy. Indeed, the relative efficacy of binding *vs.* non-binding shareholder voting is an important topic in corporate governance, both for practitioners (Allaire and Dauphin, 2016) and academics (Levit and Malenko, 2011), and my structural model is uniquely positioned to provide empirical insight.

I take my estimated structural parameters and make the SOP vote binding. In the counterfactual, the Board proposes a compensation package and shareholders hold a binding approval vote; failure leads to the CEO receiving the same compensation from the previous period. This follows the institutional design of binding votes in, for example, the United Kingdom and Switzerland (as explained in, respectively, Allaire and Dauphin, 2016; Federal Council, Switzerland, 2013).

In the model, the trade-off between non-binding and binding voting is one of discipline *vs.* information. Though a binding vote entails greater control over compensation policy, the failing pay level does not incorporate the signals revealed about CEO skill between annual shareholder meetings. This is particularly limiting early in a CEO's tenure when learning is most valuable. The binding vote's explicit control actually reduces its effective disciplining power: the shareholder optimally sets a low failure rate so that the CEO's compensation reflects current information; the Board endogenously responds to the low failure rate and the realized impact of board capture is greater.

In the counterfactual, relative to the baseline model, the SOP failure rate falls, pay levels increase and firm value decreases on average, particularly when beliefs about the CEO's ability are less precise. Indeed, the greater power of intervention actually reverses the vote's disciplining power as the Board incorporates shareholder unwillingness to fail the vote when learning is most valuable.

These findings have implications for shareholder democracy and corporate governance more generally. Levit and Malenko (2011) show that binding voting may not fully aggregate the shareholder base's information. My analysis shows that binding voting can negatively impact realized disciplining power, if failure binds the company to a policy that is not reflective of the company's current state.

The paper is organized as follows. I first describe the paper's contribution and context within the literature. Section 2 describes the data and presents empirical facts about CEO pay and SOP, which both motivate and discipline the model. Section 3 presents the structural model. Section 4 describes the estimation methodology, with Section 5 showing the results of the structural estimation. Section 6 introduces and analyzes the counterfactual binding Say-on-Pay. Appendix A contains additional results, and additional model and estimation details are in Internet Appendices IA1 and IA2, respectively.

Literature Review

This paper contributes to the literature on shareholder voice as a way to influence corporate policies (e.g. Hirschman, 1970; Gillan and Starks, 2007). Levit and Malenko (2011) study non-binding votes as a form of communication, showing how a large (activist) investor can make votes more effective at influencing management. My paper provides empirical evidence of this hypothesis by estimating how much the Board internalizes the cost of failing a SOP, and my subsample analysis shows that this cost varies with the presence of large shareholders. Levit (2019) studies the effectiveness of communication in influencing the decision-maker, which relates to my paper — the Board and shareholder costs to SOP failure determine the effectiveness of SOP as a communication device in disciplining wages.

My empirical results speak to the literature on how possibly non-binding or non-consequential shareholder voting can influence the Board of Directors. Fos and Tsoutsoura (2014) study how previous proxy contests impact the careers of directors and Aggarwal et al. (2019) study the impact of dissent votes in uncontested director elections on careers; my paper shows the career and reputation consequences of a specific form of non-binding shareholder votes — SOP. Aggarwal et al. (2023) study shareholders' motivations for voting against corporate directors and find that shareholders hold directors accountable for a wide range issues, with governance being the main driver.

My paper contributes to the literature on Say-on-Pay. Several papers study the effects of the adoption of SOP (e.g., Cai and Walkling, 2011; Ferri and Maber, 2013; Correa and Lel, 2016; Cuñat et al., 2016), showing that increasing voice through SOP improved firm value, impacted CEO pay, or both. However, given the high SOP support, several papers (such as Armstrong et al., 2013; Kaplan, 2013) conclude that, once implemented, SOP has not influenced compensation and question its effectiveness in practice. My paper shows that SOP is an effective governance mechanism: the low failure rate belies that SOP has a large impact. Relatedly, Holland et al. (2023) show that it is difficult to infer the value impact of SOP votes directly from stock prices: option-implied volatility decreases before shareholder meetings, suggesting the market internalizes the vote outcomes in advance. Dey et al. (2024) show that SOP failures trigger higher shareholder engagement on other issues, supporting my analysis that SOP acts like a performance review of the Board on a prominent policy.

There is a large literature studying how corporate governance affects firm value. Cuñat et al. (2012, 2016) and Flammer (2015) show a causal (positive) relation between adoption of governance-improving mechanisms and firm value; my paper shows how (and by how much) a particular governance mechanism improves firm value. Johnson and Swem (2021); Fos (2017) show that, although proxy contests are rare, the threat of initiation influences firm behavior beneficially for shareholders, I show that SOP operates similarly. Holderness (2018) shows that a similar mechanism exists when shareholders vote on equity issuance decisions. Within the realm of hedge fund activism, Gantchev et al. (2018) show that spillover threats of activism induce changes in corporate policies at peer firms; Zhu (2021) focuses on how (threats of) activism influence CEO compensation structures.

My paper relates to how institutional investors impact CEO compensation policy. Mehran (1995) and Hartzell and Starks (2003) show a negative relation between blockholder ownership and the level of CEO pay. Several papers have argued that passive investors, generally the largest blockholders, are ineffective monitors due to their hesitancy to dissent from management (Heath et al., 2022). In subsample analysis, I show that large blockholders are effective monitors (the Board cost to SOP failure is larger), yet they also face a larger cost to SOP failure. The argument that passive (large) investors are ineffective monitors is more subtle than previously considered, and depends on the relative magnitudes of these costs.

The study of executive compensation from an empirical, theoretical, or structural perspective is too vast to properly reference here. Taylor (2010, 2013) and Page (2018) are seminal structural papers studying CEO compensation, CEO turnover and board incentives. A structural literature that studies shareholder voting has emerged; e.g., Blonien et al. (2024) study errors in shareholder voting and Pinnington (2023) studies strategic voting considerations. To the best of my knowledge, this is the first paper to estimate a structural model of executive compensation with a shareholder vote.

Finally, SOP can be seen as a monitoring mechanism with costly punishment. While non-binding, the Board's punishment for a negative evaluation arises through a career concern or reputation channel (Dewatripont et al., 1999). Similar economic settings have been explored in the empirical industrial organization literature, for example Silveira (2017) studies how the threat of trial sentencing (and costs associated for both sides) lead to most criminal cases ending in a plea bargain.

2. Empirical Analysis

2.1. Data

For the analysis in Section 2.2 and the estimation described in Section 4, I use data on SOP vote results (Institutional Shareholder Services), CEO compensation (Execucomp) and firm-level accounting outcomes (Compustat). The period is 2011-2020 and the sample is S&P1500 firms. Table 1 displays summary statistics for firm, CEO and SOP variables. The average vote against is about 9% (i.e., the average support rate is about 91%). Only 7% of votes fail (i.e., have more than 30% vote against).

The measure of CEO compensation used throughout the paper is Execucomp's TDC1, which includes each year's salary, bonus, total value of granted restricted stock and options (the latter valued using Black-Scholes) and long-term incentive payouts. The model is silent on differences between base and incentive compensation. However, as shown below, shareholders primarily vote on the level of pay, suggesting that pay levels matter more than the structure of pay in SOP. Following arguments in Taylor (2013), I assume that each year's total pay level induces effort from the manager, and any negotiated future incentive today is incorporated into total current pay (see Page (2018) for structural analysis of the CEO's contract). The average total compensation is about 7.2 million (the median about 5.5 million), with about 30% being current compensation. The median length of tenure is 8 years.

2.2. Empirical Facts

I document empirical facts about SOP and compensation that help discipline the model. Specifically,

1. SOP failure likelihood is driven primarily by the level of total pay and company performance.
2. CEOs may exert influence over compensation policy via *board capture.*
3. SOP failure correlates with costly outcomes for directors.
4. SOP voting behavior is consistent with shareholders facing a cost from SOP failure.

Facts 1 and 2 are a summary of empirical results known to the literature, framed within my setting; Facts 3 and 4 are new results that motivate the Board and shareholder costs to SOP failure.

Fact 1. *SOP failure likelihood is driven primarily by the level of total pay and company performance.* This fact provides a basis for analysis and will inform how the SOP vote is structured in the model: the probability of SOP failure is increasing in the total pay level and decreasing in company performance. Table 2 displays regressions in which the dependent variable is an indicator for SOP failure (more than 30% voting against the SOP, in columns 1-3), or the percentage of shareholders voting against the SOP (columns 4-6). The table shows that SOP failure likelihood is strongly (and robustly) increasing in the level of total CEO pay.

The same table shows that SOP disapproval is decreasing in company performance (the firm's return on assets, or ROA), conditional on the pay level, confirming Fisch et al. (2018). These relations provide clarification for the model. Vote failure occurs when total pay is too high, given what shareholders believe about CEO ability; if the company is doing well, then shareholders are less likely to fail the SOP. Both forces will influence the shareholder voting strategy in the model.[6]

[6] As corroborative evidence that SOP impacts *future* compensation policy, Appendix Table A1 shows a strong negative relation

Fact 2. *CEOs may exert influence over their compensation via board capture.*

Previous literature has suggested that CEOs can partially determine the compensation-setting process by influencing the Board of Directors (e.g., Graham et al., 2020; Coles et al., 2014; Bebchuk and Fried, 2003). Indeed, SOP was in part designed to combat this direct channel of influence, so confirming the relation between influence on the Board of Directors and the level of compensation is important.

I examine a well-established proxy from the literature: *board co-option* (Coles et al., 2014), which measures the percentage of directors (including independent) that were appointed during the CEO's tenure. Coles et al. (2014) show that board co-option correlates with the level of CEO pay and I confirm this relation in Table A2. The level of CEO pay increases by about 8 percentages points (pp) for a standard deviation increase in board co-option.[7] Following the strong relation between board co-option and CEO pay levels, the model incorporates potential CEO influence on the Board directly into the Board's pay-setting process. In the model, the Board has a preference towards paying the CEO above the value-maximizing wage level. The magnitude of this potential bias will be controlled by a tenure-specific, estimated parameter λ_τ. As shown in Section 5.3, estimated board capture varies with empirical board co-option (the measure from Coles et al., 2014, also used in Table A2).

Fact 3. *SOP failure correlates with costly outcomes for directors.*

A basic premise of the model is that SOP failure is costly for directors. I provide new evidence of this cost in three areas. First, Table 3, columns 1-2 show that director turnover the year after the vote correlates with SOP failure. The columns show that director turnover likelihood is about 2 pp higher after SOP failure. Relative to passing, the probability of director turnover is about 20% higher. This finding is robust to controlling for company performance (ROA), and a battery of controls covering board, director and CEO features.

The table also shows a second possible cost to directors from SOP disapproval (in columns 3-4). Focusing on directors not turned over, SOP failure is associated with about a 1.3 pp higher probability of a compensation committee member leaving the committee (a 26% increase). These results show that SOP failure is a career concern for directors.

I also study how SOP failure affects the external reputation of directors, or the number of outside boards on which the director sits. Focusing on directors that sit on at least one outside board, SOP failure is associated with about a 1.8 pp increase in the likelihood that the director loses at least one outside position (a 21% increase). Vote failure seems to impact directors *outside the firm where they work.*

between changes in CEO pay and SOP disapproval: CEO pay falls by about 8 percentage points following SOP failure. This robust result suggests that SOP failure pushes the Board to make changes to CEO pay.

[7] As in Coles et al. (2014), I include CEO tenure fixed effects in each specification as co-option mechanically rises with tenure.

Fact 4. *SOP voting behavior is consistent with shareholders facing a cost from SOP failure.*

This paper presents empirical evidence that shareholders internalize a cost to SOP failure, estimates its magnitude, and explores its role in the economic mechanism through which SOP influences compensation policy. Edmans et al. (2023) provide survey evidence that of the cost: in interviews, institutional investors express their reluctance to fail SOP votes. SOP failure is viewed as a reputation cost via dissenting from management on a prominent firm policy. Shareholders also feel that dissent constitutes a future monitoring cost, as management will engage with shareholders repeatedly in the future about changing compensation policy.[8]

While the survey-based evidence suggests the existence of a shareholder cost to SOP failure, bunching in observed vote results at important failure thresholds provides empirical evidence. A higher occurrence of close passes relative to fails suggests shareholders may strategically avoid failing votes: large, pivotal blockholders may swing the outcome of the vote by (strategically) keeping the percentage of dissenting votes below the vote failure threshold. Figure A2 displays a test for bunching around the SOP failure thresholds of 30% and 50% (the commonly understood failure thresholds, ISS, 2022). The light blue and orange bars show the observed frequencies of SOP vote outcomes in 0.5 percentage point relative to the failure threshold. Though the bunching is more pronounced at the 50% threshold, it is still clearly there at 30%. Bunching is suggestive, though not definitive, evidence of a shareholder cost. While beyond the scope of this paper, a model of strategic voting could illustrate the degree to which bunching identifies a cost to failure, relative to strategic considerations (as in Pinnington, 2023).

3. Model

This section outlines the model. The key forces are guided by the analysis in Section 2. The estimation will determine the extent to which these forces matter; for now, the model treats them as parameters.

3.1. Technology and Environment

Time is annual. The firm is infinitely-lived and consists of three actors: the Board of directors, that sets the CEO's wage level each period; the CEO, who exerts effort for the wage they receive; and a shareholder base, which holds an approval vote of the Board's pay policy (a "Say-on-Pay"). As mentioned above, the wage in the model represents the level of total (expected) pay, with no distinction

[8] See Online Appendix A of Edmans et al. (2023). Investors also mention they often follow proxy advisors as resource constraints prevent them from fully analyzing compensation policy. Figure A1 presents anecdotal evidence of future monitoring costs: after failing the 2021 SOP, Netflix engaged with large shareholders about compensation numerous times throughout the year.

between base and incentive pay; one interpretation of the model is that the Board and CEO contract a new expected wage each year that induces the CEO to remain at the firm and produce (Taylor, 2013).[9]

Effort (n_t) is increasing but concave in the wage level w_t,

$$n_t = w_t^\gamma, \quad \gamma \in (0, 1].$$

This assumption captures in reduced-form that effort is privately costly for the CEO, so compensation extracts less effort if effort is already high. The firm produces according to

$$y_t = \eta A_t n_t^\beta, \quad \beta \in (0, 1),$$

where β is a constant and $A_t > 0$ is the firm's productivity. I am not interested in separately identifying γ and β, so I define $\alpha \equiv \gamma\beta \in (0, 1)$ to describe the shape of the production function.[10] The parameter η is a scaling factor, which captures how CEO productivity transformed into observed revenues.

Output is thus

$$y_t = \eta A_t w_t^\alpha. \tag{1}$$

Firm productivity is centered around CEO skill a but influenced by a mean-zero shock ε_{yt},

$$\ln A_t = a + \varepsilon_{yt}, \quad \varepsilon_{yt} \sim N\left(0, \sigma_y^2\right). \tag{2}$$

Type a is not observed by the Board and shareholders, they make predictions about a based on information they observe (as detailed in Section 3.2). Eq. (2) defines the notion of CEO skill — higher types achieve higher average productivity.

Operating income in year t is

$$\Pi_t = \eta A_t w_t^\alpha - \eta\kappa w_t.$$

$\kappa > 0$ is a model parameter which dictates how changes in CEO wages (effort) scale to firm-level costs.[11] These two parameters are not essential for the model's forces, but they are important for estimation, as they translate CEO effort and skill into the firm-level revenues and operating income observed in

[9]Effort as an increasing (concave) function of total pay can be motivated by TDC1 being primarily incentive pay, as standard principal-agent models with risk-averse agents would predict this relation. Fairness-based models like Chaigneau et al. (2022) also predict that effort may increase with pay levels, even in the absence of effort incentives.

[10]Beyond their effort and skill, the CEO is a passive actor: the model is silent on the exact contracting problem between the Board and the CEO, and instead focuses on the interaction between the Board and shareholders. Page (2018) estimates the effect of CEO attributes and agency issues on the CEO contract.

[11]$\kappa = 1$ implies costs rises in unison with output, $\kappa \lessgtr 1$ implies costs increase less or more than output.

the data (Page, 2018). Normalizing operating income by the scale parameter η (so $\pi_t = \Pi_t/\eta$), we have

$$\pi_t = A_t w_t^{\alpha} - \kappa w_t,$$

The Board of Directors (B) sets the wage w_t. Importantly, the Board does not perfectly maximize firm profits. That is, absent dynamic considerations and any influence from the SOP, the Board would choose w_t to maximize

$$\pi_t^{\mathrm{B}} = A_t w_t^{\alpha} - \kappa w_t + \lambda_\tau \kappa w_t,$$

where $\lambda_\tau \in [0, 1)$ governs the influence the CEO has on the Board's decision-making, and more generally captures agency costs in the form of CEO influence on pay, what I refer to as *board capture* (see Section 2.2 and Fact 2). I allow board capture to vary across CEO tenure according to:

$$\lambda_\tau = \lambda_0(1 + \lambda_1)^{\tau}, \tag{3}$$

where λ_0 and λ_1 are parameters to be estimated. A positive λ_τ measures a gap between the wage that would maximize shareholder value and the wage the Board would pay the CEO. Noting that output y_t is an increasing-concave function of the wage w_t, a positive λ_τ induces a realized agency cost that resembles empire-building: the Board does not internalize that higher wage levels lead to over-production (e.g., Jensen, 1986; Bebchuk and Fried, 2003).

As Taylor (2013) shows, CEO wages display *downward rigidity*: risk-averse CEOs accept lower total pay if they are protected from downside risk. To match observed patterns in compensation, the model incorporates an adjustment cost into the Board's compensation decision:

$$AC(w_t; w_{t-1}, \tau_t) = c_w \times w_t \left(\frac{w_t - w_{t-1}}{w_t} \right)^2 \times \mathbb{1}[w_t < w_{t-1}] \times \mathbb{1}[\tau_t > 0]. \tag{4}$$

The adjustment cost is quadratic, scales with the wage level, is not present in the first year of CEO tenure, and only activates if the Board decreases the wage from $t-1$ to t; the one-sidedness is chosen to match downward rigidity in CEO pay. c_w controls the cost of adjustment, and is to be estimated.[12]

Shareholders hold an approval vote each year on the Board's CEO pay policy, or a Say-on-Pay vote (SOP).[13] The vote is non-binding (failure does not lead to a change in compensation policy) and occurs

[12] Table A1 and the findings from Taylor (2013) show why the adjustment cost is needed: CEO pay falls when SOPs fail, whereas Taylor (2013) shows CEO pay is downward rigid.

[13] While the structure of the CEO's pay package certainly influences SOP outcomes, shareholders predominantly vote in

at the end of each period (after the wage decision and output have occurred, following the observed institutional structure). A failed vote results in a cost for the Board, $\chi_B \geq 0$. In practice, this cost may include both pecuniary and non-pecuniary components, but in the model it measures the Board's (perceived) aversion to a failed vote. The Board's per-period utility is

$$A_t w_t^{\alpha} - \kappa w_t + \lambda_\tau \kappa w_t - \chi_B \times \mathbb{1}[\text{SOP fail}_t] - AC(w_t; w_{t-1}, \tau_t). \tag{5}$$

Shareholders in the model seek to maximize operating income: if wages are "too high" given shareholders' current beliefs about CEO ability, the SOP vote may fail. Upon vote failure, the shareholders will also face a perceived cost, which might represent an aversion to dissenting from the Board on compensation policy (Section 2.2). The parameter $\chi_S \geq 0$ governs the shareholders' cost of failing the SOP. The shareholders' per-period utility is

$$A_t w_t^{\alpha} - \kappa w_t - \chi_S \times \mathbb{1}[\text{SOP fail}_t]. \tag{6}$$

Eqs. (5) and (6) summarize the differences in Board and shareholder preferences.[14] Shareholders can alleviate board capture (λ_τ) by threatening to fail the SOP; however, they also internalize their own cost from failed votes. The Board further internalizes the adjustment cost.

3.2. Beliefs, Signals and Model Timeline

Each period the firm separates from the CEO with exogenous probability f_τ and matches with a new CEO of tenure $\tau = 0$. Upon matching, the Board and shareholders begin with the prior about CEO skill,

$$a \sim N\left(\mu_0, \sigma_0^2\right), \tag{7}$$

which matches the distribution of ability in the CEO talent pool. At the annual compensation committee meeting, the Board receives its private signal about the CEO,

$$z_{bt} = a + \varepsilon_{bt}, \quad \varepsilon_{bt} \sim N\left(0, \sigma_{z_b}^2\right). \tag{8}$$

z_{bt} will inform the Board's wage decision, and may represent *operational interaction* with the CEO. Shareholders do not observe z_{bt}, which means that the Board is asymmetrically informed (has differ-

response to the level of CEO pay, as detailed in Fact 1 and Table 2.

[14]The assumption that shareholders do not incorporate the adjustment cost keeps the shareholders' problem static, which greatly simplifies the numerical solution.

ent and more precise beliefs) about CEO skill when they set the wage. Once the wage is set, the CEO receives their wage and exerts effort, and productivity and output realize.

At the annual shareholder meeting, each atomistic shareholder draws a signal about CEO type that is private knowledge, but correlated across shareholders. The standard voting model with incomplete information assumes that signals are completely private, but in shareholder voting, these signals are correlated (e.g., due to proxy recommendations.) This correlation makes it informationally equivalent to focus on a representative shareholder, as microfounded in Appendix IA1.1. I henceforth refer to the single representative Shareholder, labeled S. At the meeting, which occurs at the end of each period t, S aggregates information from the shareholder base into the signal

$$z_{st} = a + \varepsilon_{st}, \quad \varepsilon_{st} \sim N\left(0, \sigma^2_{z_s}\right). \tag{9}$$

Conccurently, S receives the firm's 10-K and proxy statement, which reveal output y_t, realized productivity A_t and the CEO's wage w_t. Importantly, A_t serves as a public signal about the CEO's ability: when productivity is high (2), B and S will revise their beliefs about the CEO upward. I label this signal:

$$z_{yt} = \ln A_t = a + \varepsilon_{yt}, \quad \varepsilon_{yt} \sim N\left(0, \sigma^2_y\right). \tag{10}$$

Both the private signal z_{st} and the public signal z_{yt} will affect the SOP result.

Model timeline. Figure 2 displays the per-period sequence of events. I assume signal disclosure in the wage-vote game: B's wage choice and the vote resuly fully reveal z_{bt} and z_{st}; this implies that B and S share the same beliefs $a \sim N\left(\mu_{at}, \sigma^2_{at}\right)$ at the beginning of each period t (Appendix IA1.3). At the compensation committee, B receives its signal z_{bt}, which informs their wage decision. Then, operations take place: the CEO receives their wage and expends effort, and productivity realizes and output occurs. At the annual shareholder meeting, S receives the private signal z_{st}, and output, productivity and wages are revealed; productivity reveals the private signal z_{yt}. Finally, the SOP vote occurs and Board and Shareholder utilities are realized.

Board and Shareholder beliefs. Both B and S use Bayes' rule to update beliefs about CEO ability after their signals. I use subscript a to refer to beliefs shared by B and S: $\left(\mu_{at}, \sigma^2_{at}\right)$ refers to the shared prior at the beginning t. Subscripts b and s refer to when B and S can have different beliefs. CEO tenure fully determines the variance of beliefs (see Appendix IA1 and Taylor, 2010); the function

Figure 2. Model timeline

This figure displays the within-period model timeline. The top timeline displays the timeline as it maps to practice; the bottom as it maps to the sequencing of events within each period t. Figure IA1 displays an in-depth timeline which incorporates the timing of the strategies by the Board and Shareholder (See Appendix IA1.3).

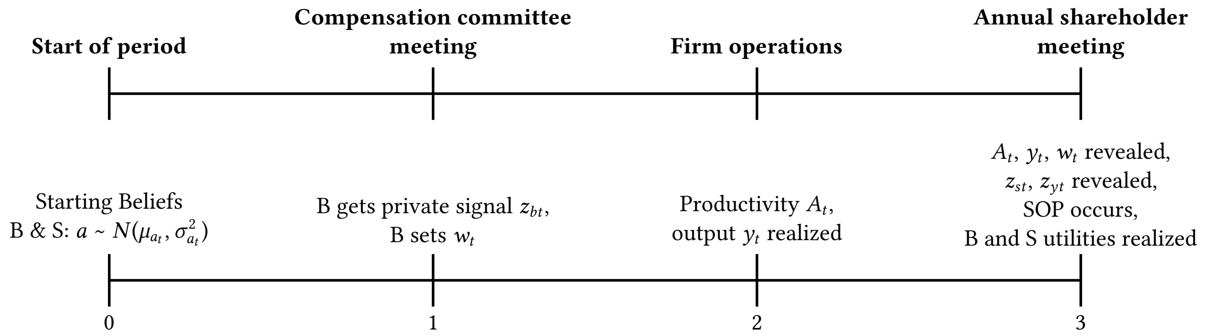

$\sigma^2_a(\tau)$ tracks how the variances of beliefs decreases across tenure. Given tenure τ_t at t,

$$\sigma^2_{at} = \sigma^2_a(\tau_t) = \sigma^2_0 \left[1 + (\tau_t) \times \sigma^2_0 \left(\sigma^{-2}_{z_b} + \sigma^{-2}_{z_s} + \sigma^{-2}_y\right)\right]^{-1}, \quad \sigma^2_a(0) = \sigma^2_0. \tag{11}$$

The mean evolves according to

$$\mu_{at+1} = \sigma^2_a(\tau_t + 1)\left[\frac{\mu_{at}}{\sigma^2_a(\tau_t)} + \frac{z_{bt}}{\sigma^2_{z_b}} + \frac{z_{st}}{\sigma^2_{z_s}} + \frac{z_{yt}}{\sigma^2_y}\right], \tag{12}$$

and the rate of decline of the variance σ^2_{at} means that μ_{at} tends toward the CEO's true ability.

3.3. The Say-On-Pay Vote

The Board and Shareholder's wage and vote decisions are informed by each party's private signal and how they view each other's beliefs. This section details the Shareholder's strategy, holding the Board's wage choice fixed. Appendix IA1.3 details three assumptions made about the SOP vote, which aim to keep the Shareholder's decision both tractable and realistic.

3.3.1. The Shareholder's Strategy

Informally, S will fail the SOP if the CEO's wage is "too high" given their beliefs. Fixing the wage choice of the Board, the notion of "too high" will incorporate S' current beliefs about CEO ability and how costly vote failure is to Shareholders. Formally, S sets a threshold posterior belief about CEO ability for which they would be indifferent between the vote failing and passing, which is equivalent to setting a threshold in S' signal distribution that leads to this posterior belief. As such, a higher threshold implies a higher probability of SOP failure. Via the Board's expected failure cost, this will lead to lower wages on average, but also raises the Shareholder's expected failure cost. S' strategy can be described as setting a probability (threat) of SOP failure which maximizes S' expected utility,

$$\Pr(\text{SOP fail}_t) = \Pr\left(Z_{syt} \leq k_{st}\right).$$

The signal Z_{syt} incorporates both the private signal (z_{st}) and the effect of firm productivity (z_{yt}) on S' beliefs. I assume the threshold S chooses takes the form

$$k_{st} = s_t \times w_t. \tag{13}$$

Fact 1: Pr(SOP fail) is increasing in the wage. Eq. (13) is chosen to match Fact 1, which shows that a higher CEO wage leads to a higher probability that the SOP vote fails. The threshold is thus increasing in the wage choice of the Board, and the choice variable s_t controls the *sensitivity* of the SOP failure likelihood to changes in the wage.

As in Figure 2, at the annual shareholder meeting, output and the wage are revealed to shareholders. S receives two signals about CEO ability: their private signal z_{st} and the productivity signal z_{yt}. The Shareholder uses all information available and considers the average of these signals, with weights determined by their relative precision. Let $p = \frac{\sigma_{z_s}^{-2}}{\sigma_{z_s}^{-2}+\sigma_y^{-2}}$ be the relative precision of ε_{st}. Hence, *ex ante* the Shareholder's signal is distributed according to

$$Z_{syt} = pz_{st} + (1-p)z_{yt} = a + p\varepsilon_{st} + (1-p)\varepsilon_{yt}, \quad Z_{syt} \sim N\left(\mu_{at}, \sigma_{at}^2 + \frac{\sigma_{z_s}^2 \sigma_y^2}{\sigma_{z_s}^2 + \sigma_y^2}\right). \tag{14}$$

This distribution determines the probability of vote failure: S incorporates both signals, placing more weight on the signal with better precision. At the time that they commit to their threshold, S has beliefs about CEO ability $\left(\mu_{at}, \sigma_{at}^2\right)$. Given this, the signal Z_{syt} has distribution as in (14). Further, B knows this is the distribution that from which S' signal arises and takes this into account when setting the wage.

Fact 1: Pr(SOP fail) is decreasing in company performance. Eq. (14) is how firm performance positively affects SOP outcomes in the model (Fisch et al., 2018). If productivity is high, S is unable to distinguish whether it is due to the CEO's expertise or a shock affecting output. A higher z_{yt} will lower the probability of SOP failure, even if the wage w_t is large.

Distance from the unbiased wage. With no board capture, the profit-maximizing, or *unbiased* wage that the Shareholder would pay, given beliefs about a, is

$$w_t^{\mathrm{U}} = \arg\max_{w_t} \quad E_{st}\left[\exp\left(a + \varepsilon_{yt}\right)\right] w_t^{\alpha} - \kappa w_t = \left(\frac{\alpha}{\kappa} E_{st}[\exp\left(a + \varepsilon_{yt}\right)]\right)^{\frac{1}{1-\alpha}}. \tag{15}$$

The Shareholder's goal is to get the Board's wage choice as close to the unbiased wage as possible, given the costliness of SOP failure. So, failure is determined via distance of the observed wage from the unbiased wage. The random variable w_t^{U} represents S' preference for the CEO wage level, given their current beliefs about CEO skill, with distribution determined by the belief tuple $\left(\mu_t, \sigma_t^2\right)$

$$w_t^{\mathrm{U}} \sim \log N\left(\frac{\mu_t}{1-\alpha} + C, \frac{\sigma_t^2 + \sigma_y^2}{(1-\alpha)^2}\right), \quad C = \frac{\log \frac{\alpha}{\kappa} + \frac{1}{2}\left(\sigma_t^2 + \sigma_y^2\right)}{1-\alpha}. \tag{16}$$

w_t^{U} is simply a transformation from belief-space to wage-space, so given the distribution of Z_{syt} in (14), there is a random variable $w_t^{\mathrm{U}}\left(Z_{syt}\right)$ given by (16) that is the conversion of Z_{syt} to its unbiased wage counterpart.[15] I refer to the CDF of this distribution as F_{st}^{U}, where st signifies the Shareholder's period t beliefs about CEO ability.[16]

3.3.2. Determining the Probability of Say-on-Pay Failure

If Z_{syt} leads to S' posterior belief falling below the chosen the threshold, the vote fails. Given the previous discussion, the *ex ante* probability of failure, given w_t is

$$\Pr(\text{SOP fail}_t) = \Pr\left(w_t^{\mathrm{U}}\left(Z_{syt}\right) \leq s_t \times w_t\right) = F_{st}^{\mathrm{U}}(s_t \times w_t). \tag{17}$$

Fixing the Board's best response w_t for now, S chooses s_t to maximize expected operating income, conditional on their beliefs at the start of period t, $a \sim \left(\mu_{at}, \sigma_{at}^2\right)$ and the signals they will receive at the end of the period. Importantly, S influences expected wages. When setting the vote policy, S takes expectations over signals z_{bt}. Concurrently, B updates beliefs and optimally offers $w_t(z_{bt}, s_t)$. The Shareholder's problem is

$$\max_{s_t} \quad \int_{z_b} f(z_b) \left[\underbrace{E_{st}[A_t] w_t(z_{bt}, s_t)^{\alpha} - \kappa w_t(z_{bt}, s_t)}_{\text{Expected operating income}} \quad \underbrace{- \chi_S \times F_{st}^{\mathrm{U}}(s_t \times w_t(z_{bt}, s_t))}_{\text{Expected cost of SOP failure}} \right] dz_b. \tag{18}$$

[15]The distribution of $w_t^{\mathrm{U}}\left(Z_{syt}\right)$ is found by plugging μ_{at} and $\sigma_{at}^2 + \frac{\sigma_{z_s}^2 \sigma_y^2}{\sigma_{z_s}^2 + \sigma_y^2}$ into (16).

[16]The transformation to the lognormal wage distribution ensures that S' threshold is always increasing in the wage. While Z_{syt} is normal and has positive and negative support, the lognormal wage is guaranteed to be non-negative.

The Shareholder commits to an *ex ante* threat of vote failure, and the realizations of z_{bt} (determining the wage), z_{st} and z_{yt} will determine the outcome of the vote.

Commitment to a costly failed vote. It is important to note that I assume Shareholder commitment to an *ex post* costly failed vote. This may seem a strong assumption but reflects the nature of a costly punishment mechanism. If the Shareholder reneged on the failed vote (after inducing the Board to pay a more favorable wage), this would affect the Shareholder's credibility and effectively lower the vote's disciplining power. Indeed, as Fact 3 shows, director turnover increases after failed votes, suggesting a willingness to punish the Board, though it may entail a cost for the Shareholder.

3.4. The Compensation Committee

Each period, the Board receives its signal z_{bt} and decides the CEO's wage. Their beliefs at the beginning of t are $a \sim N(\mu_{at}, \sigma^2_{at})$. Upon receiving z_{bt}, the Board updates to $\left(\mu_{bt|z_b}, \sigma^2_{bt|z_b}\right)$, where

$$\begin{aligned}\mu_{bt|z_b} &= \sigma^2_{bt|z_b}\left(\frac{\mu_{at}}{\sigma^2_{at}} + \frac{z_{bt}}{\sigma^2_{z_b}}\right) \\ \sigma^2_{bt|z_b} &= \left(\sigma^{-2}_{at} + \sigma^{-2}_{z_b}\right)^{-1} = \frac{\sigma^2_{at}\sigma^2_{z_b}}{\sigma^2_{at} + \sigma^2_{z_b}}.\end{aligned} \tag{19}$$

If B's belief is revised downwards, they will want to decrease w_t relative to w_{t-1}, so I include the adjustment cost from (4) in the Board's problem (to match rigidity in total pay, Taylor, 2013). At the compensation committee meeting, B's wage policy solves the following Bellman equation:[17]

$$\begin{aligned}V(\mu_{at}, \tau_t, w_{t-1}) = \max_{w_t} \quad & E_{bt|zb}[A_t]w_t^{\alpha} - \kappa w_t + \lambda_\tau \kappa w_t - \chi_B \times F^{\mathrm{U}}_{st}(s_t \times w_t) - AC(w_t, w_{t-1}; \tau_t) + \\ & \delta_B\Big[(1 - f_{\tau_t})E_{bt|zb}[V(\mu_{at+1}, \tau_t + 1, w_t)] + f_{\tau_t} V^R\Big].\end{aligned} \tag{20}$$

The state consists of the two variables that track the Board's beliefs: the current belief about the mean μ_{at} and the CEO's tenure τ_t (which determines the variance of beliefs, eq. 11). The third is the previous period's wage w_{t-1} (set to zero if $\tau_t = 0$). $F^{\mathrm{U}}_{st}(s_t \times w_t)$ specifies the probability of SOP failure as detailed in Section 3.3. Operator $E_{bt|zb}[\cdot]$ is taken with respect to B's beliefs about CEO ability after it receives the signal z_{bt} (so beliefs are distributed according to eq. 19).

The Board faces a trade-off between paying a higher wage (λ_τ) and the increased probability of SOP failure. It must also consider the firm's expected productivity A_t (influenced by ε_{yt}), the likelihood of

[17] w_t is a function of both z_{bt} and the shareholder's strategy s_t, however for tractability this notation is omitted from (20).

SOP failure, adjustment costs, and the continuation value. The hazard function f_{τ_t} controls CEO-firm separation and is an input to the model.[18] V^R describes the termination value: upon separation, the Board finds a new CEO and beliefs reset to (μ_0, σ_0),

$$V^R = V(\mu_0, \tau = 0, 0). \tag{21}$$

3.4.1. Objective Firm Value

The Board's optimal wage policy admits the model's definition of firm value. Given the state and the Board's wage policy $\omega_t = \omega(\mu_{at}, \tau_t, w_{t-1})$, firm value is simply the discounted cash flows of the firm:

$$V^{\text{OBJ}}\Big(\mu_{at}, \tau_t, w_{t-1}\Big) = E_{at}[A_t]\omega_t^{\alpha} - \kappa\omega_t + \delta_B \times E_{at}\Big[V^{\text{OBJ}}\Big(\mu_{at+1}, \tau_t + 1, \omega_t\Big)\Big]. \tag{22}$$

Unlike the Board's problem, objective value does not include board capture or the adjustment cost; in general, the Board's policy ω_t will be above the value-maximizing wage. Secondly, firm value is not directly affected the SOP vote. In counterfactuals, (22) will allow me to analyze how changes to SOP affect firm value (via its effect on the Board's compensation policy).

3.5. Model Solution

I use Bayes' rule to derive Board and shareholder beliefs about CEO ability and substitute these beliefs into the Board's Bellman equation (20) and the Shareholder's objective function (18). To solve, I guess the Shareholder's solution for s_t, given each Board's state (μ_t, τ_t, w_{t-1}). Given S' proposed vote policy, I iterate to find the Board's policy function $\omega(\mu_t, \tau_t, w_{t-1})$, at which point I update S' vote policy. I continue until both the wage and vote policies are stable. Appendix IA1.4 gives the full solution.

4. Estimation

I estimate the parameters of the structural model using indirect inference (McFadden, 1989): I choose the vector of structural parameters which minimizes the difference between the reduced-form outcomes of an auxiliary model estimated on observed and simulated data. The auxiliary model, though misspecified, focuses on features of the data which are highly informative about structural parameters. Details of the estimation are presented in Appendix IA2. Section 4.1 outlines the identification strategy and Section 4.2 presents the quality of the model's fit.

[18]Pooling all CEO spells, f_{τ_t} is the frequency of turnover after τ years, conditional on the CEO surviving $\tau - 1$ years. I follow Taylor (2010) in the computation of the hazard rates. For simplicity, $f_0 = 0$ in the estimation, and $f_{\tau_t} = 1$, where **T** is the cap on the length of CEO tenure. Further, the estimation sample excludes CEO spells that only last one year.

4.1. Identification Strategy

Output and CEO skill parameters. I first residualize out industry and time fixed effects from log firm revenues.[19] The model's definition of log revenues is:

$$\log y_{it} = \log \eta + a_i + \alpha \log w_{it} + \varepsilon_{yit} \implies y_{it} = \eta A_{it} w_{it}^{\alpha} = \eta \exp\left(a_i + \varepsilon_{yit}\right) w_{it}^{\alpha}. \tag{23}$$

y_{it} is CEO-driven firm revenues, w_{it} is the CEO's observed compensation, and η scales CEO productivity and effort up to observed levels. I target the parameter $\log \eta$ in the estimation. Average CEO skill is only identified relative to the constant $\log \eta$ in (23), so I normalize μ_0 to zero. The following regression in the data maps exactly to company output in the model

$$\log y_{it} = y_0 + y_1 \log w_{it} + \epsilon_{it}^{y}. \tag{24}$$

Given $\mu_0 = 0$, $\hat{y}_0$ (average observed revenues) identifies $\log \eta$. The curvature of output with respect to the wage/effort (α) is identified via $\hat{y}_1$, and σ_y via the variance of the residual $\widehat{Var}\left(\epsilon_{it}^{y}\right)$. Netting out the effect of CEO effort/wage on output and taking expectation across the years CEO i spent in office exactly pins down σ_0 (Taylor, 2010):

$$E_i[\tilde{y}_{it}] = E_i[\log y_{it} - \alpha \log w_{it}] = E_i\left[a_i + \varepsilon_{yit}\right] = a_i \implies Var(E_i[\tilde{y}_{it}]) = \sigma_0^2. \tag{25}$$

Lastly, targeting average observed profit margins identifies the scale parameter κ. Unscaled operating profits from the model are $\Pi_t = \eta A_t w_t^{\alpha} - \kappa \eta w_t$, so

$$\frac{\Pi_{it}}{y_{it}} = \frac{A_t w_t^{\alpha} - \kappa w_t}{A_t w_t^{\alpha}}. \tag{26}$$

Parameters that drive the Board's decision. The optimal wage depends on the Board's beliefs about CEO productivity, the degree of board capture, the probability of SOP failure, and an adjustment cost if the Board lowers pay. These forces are reflected in the following approximation to the optimal log wage, illustrating how estimable parameters drive variation in CEO pay

$$\log w_{it} \approx \underbrace{\log(\alpha E_{bt}[A_t])}_{\text{Expected productivity}} - \underbrace{\log \kappa(1 - \lambda_\tau)}_{\text{Board capture}}$$

[19] That is, I regress log firm revenues on industry and time fixed effects (both assumed mean zero) and take the residual as my measure of firm revenues.

$$- \underbrace{\log\left(1 + \frac{\chi_B}{\kappa(1-\lambda_\tau)} \times f^{\text{U}}_{sit}(s_{it} \times w_{it})s_{it}\right)}_{\text{Pr(SOP fail)}} \quad + \quad \underbrace{g(w_{it}, w_{it-1})}_{\text{Adjustment cost}} ,$$

where $f^{\text{U}}_{sit}(\cdot)$ is the likelihood function of the Shareholder's signal distribution.[20]

The following regression in the data is tightly linked to this expression:

$$\log w_{it} = b_0 + b_1 \mathbb{1}[\text{SOP fail}_{it}] + b_2 \tau_{it} + \phi_i + \epsilon^b_{it}. \tag{27}$$

where ϕ_i is a CEO fixed effect and τ_{it} is the CEO's tenure. The fixed effect ϕ_i will remove variation caused by natural CEO ability (the expected productivity component of the wage); so $\hat{b}_0$ and $\hat{b}_2$ will capture λ_0 and λ_1, respectively. $\hat{b}_1$ helps to identify χ_B as it indicates how different wages are in SOP failure relative to SOP pass. Fixing other parameters, a small χ_B (close to zero) implies that b_1 must be close to zero: the Board cares little about SOP and will set a similar wage in SOP pass and failure. As χ_B gets larger, the Board must have better beliefs about the CEO to offset the higher expected cost of failure, resulting in a larger b_1. The Board's wage choice is determined by their signal z_{ibt}, so $\widehat{Var}\left(\epsilon^b_{it}\right)$ informs about the precision of the Board's signal. Lastly, the log wage autocovariance

$$\gamma(w_{it}, w_{it-1}) = Cov(\log w_{it}, \log w_{it-1}) \tag{28}$$

identifies c_w, as it pins down the dynamic aspect of the Board's problem.

Parameters that drive the Shareholder's decision. The model's definition of SOP failure is

$$\mathbb{1}[\text{SOP fail}_{it}] = \mathbb{1}\left[s_{it} w_{it} - w^{\text{U}}_t\left(p z_{sit} + (1-p) z_{yit}\right) \geq 0\right],$$

where $w^{\text{U}}_t\left(p z_{sit} + (1-p) z_{yit}\right)$ is a linear combination of the two signals z_{sit} and z_{yit} transformed to the relevant distribution as in (14), and s_{it} is the Shareholder's choice variable. The model-implied determination of SOP failure maps to the regression in the data

$$\mathbb{1}[\text{SOP fail}_{it}] = s_0 + s_1 \log w_{it} + s_2 \epsilon^y_{it} + +\phi_i + \epsilon^s_{it}, \tag{29}$$

[20]The expression for $\log w_{it}$ results from taking w_{it-1} as given, fixing the Shareholder's strategy s_{it} and abstracting from dynamics (setting $\delta_B = 0$). The Board's first-order condition from this simple program is the approximation. The partial derivative of the adjustment cost with respect to w_t can be written as as $AC'(w_t, w_{t-1}) = c_w\left(1 - \frac{w^2_{t-1}}{w^2_t}\right)\mathbb{1}[w_t < w_{t-1}]\mathbb{1}[\tau > 0]$. Then, $g(w_t, w_{t-1}) \approx \log\left(2 + \frac{AC'(w_t, w_{t-1})}{1+\chi_B \times (\kappa(1-\lambda_\tau))^{-1} s_{it} f(s_{it} w_t)}\right)$.

where ϵ_{it}^{y} is the residual from the output regression (24). $\hat{s}_0$ maps to the unconditional failure rate and is most informative for the parameters χ_S and λ_τ. All else equal, higher χ_S will lower the observed rate of SOP failure and higher λ_τ will raise it (via increasing the wage). $\hat{s}_1$ describes the sensitivity of SOP failure likelihood to the wage, which is highly informative about χ_S and χ_B: s_1 is decreasing in χ_S as it lowers the sensitivity of the threat of SOP failure to the wage and increasing in χ_B as the Shareholder internalizes the threat has a larger impact on wages.

While the Board's wage choice reveals z_{bit}, the Shareholder's signal is unobservable in the data. As the distribution of ϵ_{it}^{y} is determined by other parameters (σ_y and σ_0), $\widehat{Var}(\epsilon_{it}^{s})$ and $\hat{s}_2$ jointly identify σ_{z_s}. $\hat{s}_2$ corners how changes in productivity influence Shareholder beliefs about the CEO and determines how much of SOP is driven by the private signal z_{sit}.

4.2. Estimated Model Fit

4.2.1. Moment-Matching Exercise

Table 4 displays the closeness of auxiliary-model moments estimated on the observed and simulated data. The final column displays the test statistic from a two-way t-test comparing each moment. Overall, the fit is quite good and importantly, the estimation matches the reduced-form moments which identify key structural parameters λ_τ, χ_B and χ_S.

The estimation matches the observed SOP failure rate closely (moment 1, $\hat{s}_0$): 6.8% and 6.4% in the observed and simulated data, respectively. The sensitivity of SOP failure to the wage (moment 2, $\hat{s}_1$) and the output residual (moment 3, $\hat{s}_2$) are well-matched. These three moments are key for identifying the Shareholder's voting preference, so the closeness is reassuring. The vote regression residual variance is slightly too large in the simulated data (moment 4), and the statistically significant difference primarily reflects the high precision of this moment in the data.

For the Board, average log wages (moment 5, $\hat{b}_0$) and the log difference in wages in SOP failure (moment 6, $\hat{b}_1$) are matched well; the model under-estimates the average wage when SOP votes pass by about $24 thousand, and the average wage when SOP votes fail by about $159 thousand. The tenure time trend in wages (moment 7, $\hat{b}_2$) is the the primary area in which the model fails to replicate the data: the model has too-low of a within-CEO wage growth to match the data. This suggests that a high growth of board capture is needed to keep wages increasing over tenure as the disciplining power of SOP grows as beliefs become more precise. Otherwise put, a relatively high degree of board capture is needed to keep wages from decreasing over tenure. The variance of the wage regression residual and the log wage autocovariance (moments 8 and 9) are indistinguishable in the observed and simulated

Figure 3. Untargeted moment: The relation between changes in CEO pay and SOP vote outcomes
This figure displays the relation between changes in CEO pay and the SOP vote outcome in both the observed and simulated data. In both panels, I display a binned scatterplot estimated from a regression of the change in CEO pay on SOP disapproval, expressed as the vote's distance from the failure threshold. Each regression includes CEO and CEO tenure fixed effects, and a control for the output shock from the previous period. Both the change in compensation and the SOP vote results (expressed as distance from the vote failure threshold) are de-meaned in the plot.

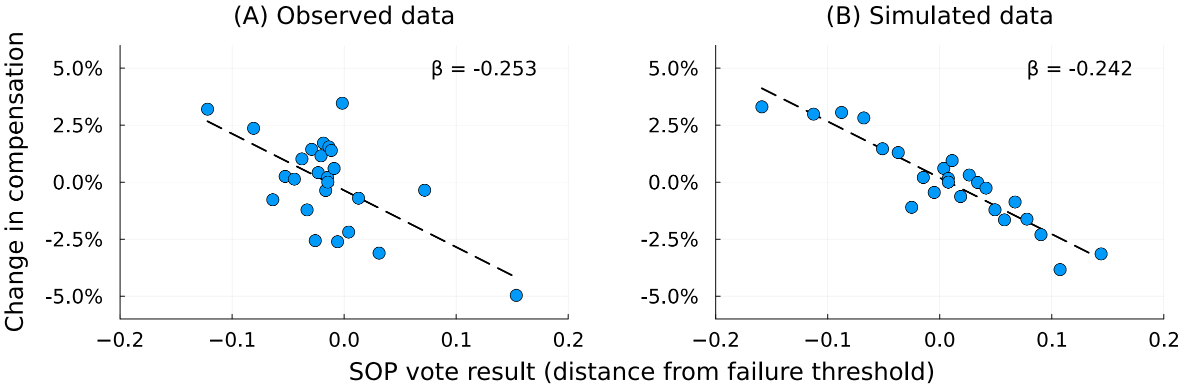

data, suggesting that the Board's private information volatility and adjustment cost are well-identified.

Output and skill-based moments are matched well in magnitudes (moments 10-14). The output residual variance is close in magnitude, but given the precision of this moment in the data, the difference is statistically different from zero. The model also estimates profit margins to be too high.

4.2.2. Changes in CEO Pay Following Say-on-Pay Disapproval

Confidence in the model's implications will be strengthened if the simulated and observed data produce similar results when examining features of the data not target during estimation. Figure 3 displays how the observed and simulated data compare in terms of *changes* in CEO compensation (from t to $t + 1$) in response to different SOP disapproval rates (in t). The slope of this relation is very similar in the model and data (though understandable more precise in the model). A key prediction of the model is that SOP disapproval conveys that shareholders believe the CEO is low-type. The figure shows that the Board endogenously internalizes shareholder beliefs into future pay decisions.

5. Results

5.1. Estimated Structural Parameters and Economic Implications

Table 5 displays the estimated parameters and quantifies their economic magnitudes.

5.1.1. Board Capture

The estimated values for λ_0 and λ_1 are 0.131 and 0.066, respectively. In Appendix IA1.5, I show that the share of total surplus shifted to the CEO by board capture can be written as

$$\theta(\tau) = \frac{\lambda_\tau w_\tau}{E_\tau[\pi_\tau] + \lambda_\tau w_\tau} = \frac{\lambda_\tau w_\tau}{E_\tau[A]w_\tau^\alpha - w_\tau + \lambda_\tau w_\tau}. \tag{30}$$

$\theta(\tau)$ would equal zero without board capture and would equal one if the CEO captured all profits, and importantly it incorporates how the Board's preferences for w_τ change as it learns about CEO ability. Figure 4 displays estimated CEO surplus capture as a function of tenure. Focusing on a CEO of average skill in their first year of tenure, I estimate $\theta(0) = 16.4\%$, which rises to 28.6% and 58.1% in the 10th and 20th years of tenure, respectively. This is similar to Taylor (2013), who estimates that CEOs capture around 50% of surplus on the upside, but bear no downside risk. I allow my measure of CEO surplus capture to vary across tenure, following prior empirical research (e.g., Coles et al., 2014). These estimates suggest that CEO surplus capture is substantial at public companies.

5.1.2. The Board and Shareholder Costs to Say-on-Pay Failure

I estimate that the Board considers SOP failure to be equivalent to 3.14% of firm value. For shareholders, this cost is 0.90% of value (or of a shareholder's equity stake), and both estimates are statistically different from zero.[21] It is important to note that these are utility costs: SOP failures do not affect value directly, the Board and shareholders must behave *as if* they do for the model to match observed outcomes. Thus, the shareholder cost to SOP failure makes shareholders pass some SOP votes that would fail if this cost did not exist. Further, both costs are considerably lower in expectation as the unconditional failure rate is 6.4% in the simulated data: the Board (Shareholder) expected cost is closer to 0.21% (0.05%) of value. These estimates highlight a key finding of this paper and clarify the economic mechanism of Say-on-Pay votes. SOP resembles a costly punishment mechanism: shareholders can punish the Board for overpayment, but internalize a cost from doing so.[22] The threat of costly SOP failure disciplines the Board even when failure is unlikely. As I show in Section 5.2 providing shareholders with a punishment technology improves firm value considerably, *even though* punishing the Board is conditionally costly.

[21] The parameters capturing the failure cost to the Board χ_B and Shareholder χ_S are estimated to be 2.247 and 0.641. I normalize them by average firm value (in Appendix IA1.5, I show that average firm value can be derived in closed form).

[22] Costly punishment is a common, naturally occurring mechanism that facilitates cooperation between economic actors, and has been analyzed by a large experimental literature (e.g., Ambrus and Greiner, 2012).

Figure 4. Estimated CEO surplus capture over CEO tenure

This figure displays estimated CEO surplus capture from (30) over the first 20 years of CEO tenure. Board capture λ_τ is estimated using λ_0 and λ_1 from Table 5.

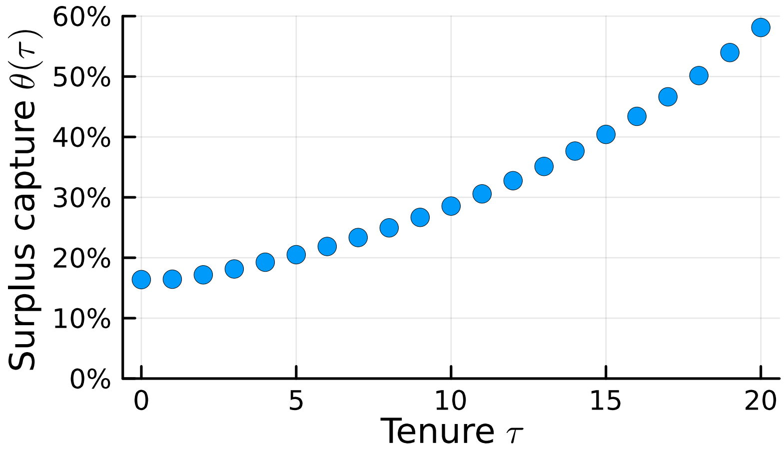

5.1.3. Economic Implications of Other Parameters

The volatility of CEO skill is 1.105, equating to about 1.5% of average firm value. This implies that CEO skill matters (supporting Taylor (2010), who finds that CEO fixed effects matter greatly). The output shock volatility σ_y is 1.465 (about 2.1% of firm value): a large portion of observed variation in output comes from randomness, rather than CEO skill.

The standard deviations of the Board's and shareholder's private signals are 0.818 and 0.684, respectively. While the volatility of the Board's innovation is larger than the Shareholder's, the Board receives its signal in advance of making its wage choice; at the point when B and S set their strategies, the Board has more precise beliefs than the Shareholder. Figure A3 displays how the variance of (Board and Shareholder) beliefs decline over the CEO's tenure. There is still quite a lot of uncertainty through the median length of the CEO's tenure ($\tau = 8$, as in Table 1).

5.2. How Much Does Say-on-Pay Impact Compensation and Firm Value?

The Board SOP failure cost suggests that SOP impacts compensation policy, but the estimated magnitude does not reveal this impact directly. Setting $\chi_B = 0$, simulating a counterfactual dataset and comparing quantities reveals the full impact of SOP, and also allows me to benchmark my estimation against empirical work on how the *adoption* of SOP impacted compensation policy and firm value.

Table 6 displays the results. As this counterfactual is equivalent to removing SOP, I do not display the SOP failure rate. The table displays the percentage that pay levels would increase if SOP were removed, both for the full simulated sample and for several sample splits. It shows that, on average, SOP brings pay down by 6.6%. This is consistent with the evidence from Correa and Lel (2016) who show (in a cross-country analysis) that total CEO pay decreased by 7% upon the adoption of SOP; the similarity in our estimates is reassuring.

There is cross-sectional heterogeneity in the impact. SOP has more impact for below-average ability CEOs (columns 2-3, 8.5% decrease *vs.* 4.8%), and for CEOs later in tenure (columns 4-5, 10.3% *vs.* 3.6%). This suggests that SOP has more impact in the cases when it should: board capture negatively impacts performance more for low-ability and longer-tenured CEOs (board capture increases in tenure). Interestingly, the change in the wage level is greatest for the observations in which SOP fails (column 6, 13.8% decrease): the cases when discipline actually occurs are when SOP is most impactful.

The table also shows that SOP increases firm value by 2.35% on average, with similar heterogeneity as observed in compensation. This impact on value occurs because changes in compensation bring the CEO's effort closer to that which maximize shareholder value. Cuñat et al. (2016) find that the adoption of SOP increased market value by 5%. Their design uses firms voting to adopt SOP (before it was mandated by Dodd-Frank), and is based on the discontinuity at the majority threshold in the adoption vote. Given their estimate is local to the firms close to the threshold (and thus likely to benefit more than the average firm), it is reassuring that my average estimate is somewhat smaller.

This analysis benchmarks the key outcomes of the estimation against two important papers that study how the adoption of SOP affected compensation policy and value. It also shows that shareholder voice via a regularly-occurring opportunity to dissent from the Board on compensation policy benefits shareholders. The threat of punishment is value-enhancing, though punishments rarely occur.

5.3. Subsample Analysis

I use subsample estimation to explore how parameters and their implications vary with empirical measures of board capture (board co-option, Coles et al., 2014) and institutional ownership concentration (percentage of market value held by the 5 largest investors, Hartzell and Starks, 2003). I estimate the model on each subsample and display the resulting parameters from each split in Table 7.[23]

5.3.1. Board Co-Option

My estimation shows that CEOs receive a large share of surplus. While the model takes no stance on the channel through which board capture arises, Table A2 shows that *board co-option* (Coles et al., 2014), which measures the percentage of the Board appointed during the tenure of the CEO, may play a role:

[23]Table IA2 displays the parameters (with standard errors) and moment-matching exercise for each subsample.

if CEOs influence director selection to tilt Board decisions in their favor, then estimated board capture should vary with empirical board co-option. I split the sample by the degree of board co-option. Board co-option mechanically increases with tenure, so I use tenure-residualized board capture, averaged across the CEO's tenure at the firm, as the splitting variable. High (low) co-option firms are in the top (bottom) quintile of board co-option.

Table 7 columns 1-3 display the results of this subsample estimation. For the low co-option sample, the table shows that $\lambda_0 = 0.126$ and $\lambda_1 = 0.082$; and for the high co-option sample, $\lambda_0 = 0.099$ and $\lambda_1 = 0.105$. Although the estimated baseline board capture (λ_0) is higher in the low co-option sample, CEO surplus capture increases at a faster rate in the high co-option sample, averaging about a 3 percentage point greater growth per year compared to the low co-option group. Interestingly, as Table 7 shows, the board failure cost is much larger in the low co-option (4.28% *vs.* 1.33%), while the Shareholder costs are practically equivalent in the two samples. This suggests that board capture may reduce the effectiveness of SOP as a governance mechanism.

5.3.2. Institutional Ownership Concentration

Institutional owners, by their size and consequent influence on corporate policies, may take on the role of disciplining management (Kakhbod et al., 2023; Appel et al., 2016; Brav et al., 2008). Hartzell and Starks (2003) find that institutional ownership concentration is negatively related to levels of CEO pay: the presence of large blockholders disciplines compensation policy. However, recent research argues that the largest blockholders (passive funds) may be ineffective monitors (Heath et al., 2022), as they tend to vote with management more regularly.

If large blockholders discipline compensation policy, the model predicts that the Board's SOP failure cost should be higher with their presence. The model also implies that the Shareholder SOP failure cost should increase with blockholder concentration. When the shareholder base is dispersed, no single investor may be focal enough to deal with the potential fallout of a failed vote (see Fact 4 and Edmans et al., 2023). So, the model predicts that the shareholder cost to SOP failure should also increase with blockholder concentration.

I follow Hartzell and Starks (2003) and use the percentage of the firm's market capitalization held by the top five institutional investors (top 5 inst. ownership) as a measure of concentration. Table 7 display the estimated parameters for the sample split into "low" and "high" based on the median average top 5 inst. ownership over the CEO's tenure. The Board's failure cost is higher when there is greater concentration of institutional investors. Table 7 columns 4-5 (row 1) show that the Board cost is 3.62%

(4.48%) of firm value for the low (high) split based on top 5 inst. ownership, and the difference is statistically significant. The Shareholder cost also increases with concentration, going from 1% to 2.66%. The estimation reveals that large blockholders do discipline compensation policy, as the punishment they can inflict on the Board is larger. However, the cost of *giving* this punishment is also larger.

6. Counterfactual: Should Say-on-Pay Votes Be Binding?

My analysis so far has taken as given the way Say-on-Pay occurs in practice. Though non-binding vote failure is conceivably ignorable by boards, my results suggest that SOP emerges as an effective governance mechanism. What is not clear is whether SOP would be more effective as a *binding* vote. Indeed, the relative efficacy of binding vs. non-binding shareholder voting has long been discussed by academics and practitioners, both concerning SOP (e.g., Ferri, 2015; Allaire and Dauphin, 2016) and more generally (e.g., Levit and Malenko, 2011; Levit, 2020; Kakhbod et al., 2023). Despite this, there is scant empirical evidence comparing binding *vs.* non-binding voting. Such analysis would require comparison of outcomes (with respect to the policy and firm value) for the same firm under a binding and non-binding regime, and such a setting is unlikely to exist in the data.

My structural model is thus uniquely positioned — I can alter the structure of SOP and ask the question: is a binding SOP vote a more valuable governance mechanism? The answer has important implications for corporate governance and shareholder democracy, beyond this particular setting.

In the baseline model, SOP is purely shareholder voice: the Shareholder votes on realized compensation and the only consequence of a failed SOP vote is an *ex post* utility cost. Under the binding counterfactual, vote failure would lead to the CEO receiving the previously approved remuneration package (i.e., last period's compensation). This resembles how binding SOP votes work in the UK (these votes occur triennially in concert with annual non-binding SOP votes); rejection of compensation policy in the UK's binding vote requires the company to operate according to previous compensation policy (Allaire and Dauphin, 2016).[24]

Figure 5 provides a view of the institutional and model differences between a non-binding and binding vote. In practice, non-binding SOP occurs *ex post* and is simply an approval vote at the end of the fiscal year. In the counterfactual, SOP must be *ex ante* and occur after the Board has proposed compensation policy but before the CEO is paid. As such, vote failure would imply a different compensation package for the CEO and consequently, different company output and performance.

[24] Further, a similar policy exists in Switzerland, by the Ordinance Against Excessive Remuneration, shareholders must approval the total compensation of top executives, and should shareholders reject proposed pay, the CEO continues to receive compensation according to the last approved package (see Federal Council, Switzerland, 2013).

Figure 5. Binding SOP: Counterfactual model timeline

This figure displays the within-period model timeline of the counterfactual, in which SOP is a binding vote. The top timeline ("Baseline") displays the timeline from the main version of the model, the bottom timeline ("Counterfactual") displays the counterfactual timeline.

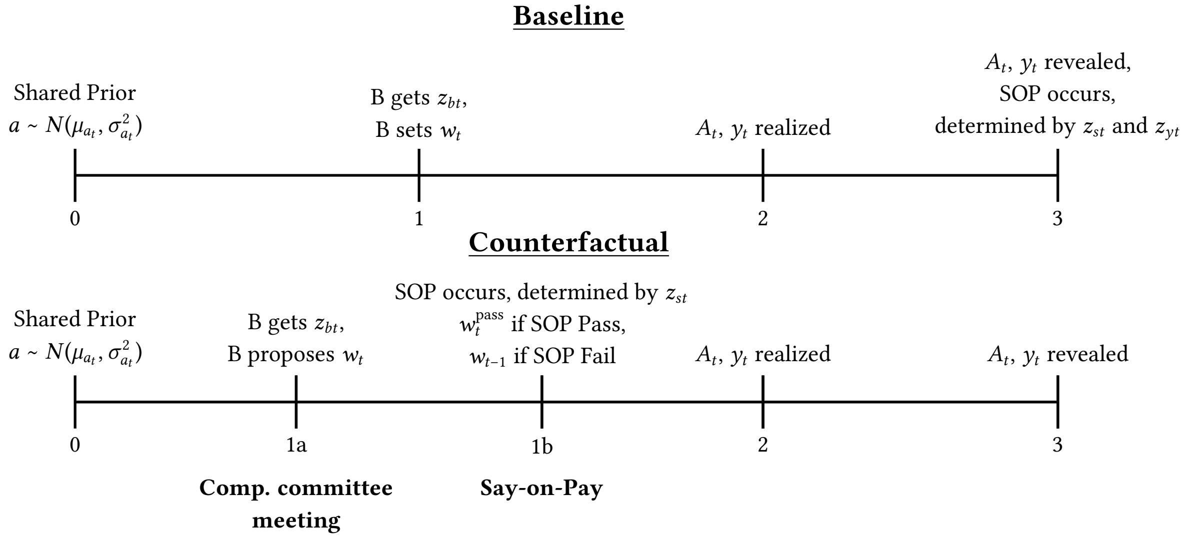

6.1. The Compensation Committee and Binding Vote

As in Section 3.4, the Board and Shareholder come into period t with the shared prior about CEO ability $a \sim N\left(\mu_{at}, \sigma^2_{at}\right)$. At the compensation committee meeting, instead of setting a finalized remuneration package, B only *proposes* a wage level at the compensation committee. During the SOP vote, the Shareholder votes on the Board's proposed pay. The counterfactual pay level takes the form

$$w_t = \begin{cases} w_t^{\text{pass}} & = \tilde{\omega}(\mu_{at}, \sigma^2_{at}, w_{t-1}) \qquad \text{if SOP pass} \\ w_t^{\text{fail}} & = w_{t-1} \qquad \text{if SOP fail} \end{cases} . \tag{31}$$

When the vote passes, the CEO receives the Board's proposed wage (determined by policy function $\tilde{\omega}(\cdot)$, detailed below), set according to Board beliefs after they receive their signal z_{bt}, as in the baseline. When the vote fails, the CEO receives the same pay as the previous period, w_{t-1}. The binding vote entails greater discipline for the Shareholder at the expense of a compensation package that does not fully internalize all information about the CEO's ability: w_{t-1} does not incorporate the signals revealed about CEO ability between compensation committee meetings (namely, the previous period's output and Shareholder signals z_{yt-1} and z_{st-1}, and the Board's current-period signal z_{bt}).

The Shareholder's problem. The Shareholder's problem is markedly different. First, S now has an active role in compensation policy. Second, the information that informs the SOP vote differs in the binding vote — information coming from productivity (z_{yt}) no longer informs the SOP vote (only the private signal z_{st} has an impact). For notational convenience, I define

$$F^{\text{U}}_{st}(s_t \times \text{w}) = \Pr(\text{SOP fail} \mid s_t \times \text{w})$$
$$\tilde{\pi}_{st}(\text{w}) = E_{st}[A_t]\text{w}^{\alpha} - \kappa\text{w}.$$

as SOP failure likelihood (which follows the same definition as eq. 17) and the firm's operating income, respectively, given wage w. S controls the sensitivity of SOP failure to the proposed wage, w_t^{pass}

$$\max_{s_t} \quad \left(1 - F^{\text{U}}_{st}\left(s_t \times w_t^{\text{pass}}\right)\right) \times \underbrace{\int_{z_b} f(z_b)\left[\ \tilde{\pi}_{st}\left(w_t^{\text{pass}}\right)\right]\, dz_b}_{\text{Non-binding: CEO receives proposed wage}} +$$
$$F^{\text{U}}_{st}\left(s_t \times w_t^{\text{pass}}\right) \times \underbrace{\left(\tilde{\pi}_{st}(w_{t-1}) - \chi_S\right).}_{\text{Binding: CEO receives prior wage \& S pays cost } \chi_S} \tag{32}$$

The Board's problem. Under binding SOP, the Board proposes w_t^{pass} according to

$$\tilde{V}(\mu_{at}, \tau_t, w_{t-1}) = \max_{w_t^{\text{pass}}}\left(1 - F^{\text{U}}_{st}\left(s_t \times w_t^{\text{pass}}\right)\right) \underbrace{\left[(\tilde{\pi}^{\text{B}}_{at|z_{bt}}(w_t^{\text{pass}}) - AC(w_t^{\text{pass}}, w_{t-1}; \tau_t) + \tilde{V}_{at|z_{bt}}(w_t^{\text{pass}})\right]}_{\text{Non-binding: Proposed compensation}} +$$
$$F^{\text{U}}_{st}\left(s_t \times w_t^{\text{pass}}\right) \underbrace{\left[\tilde{\pi}^{\text{B}}_{at|z_{bt}}(w_{t-1}) + \tilde{V}_{at|z_{bt}}(w_{t-1}) - \chi_B\right]}_{\text{Binding: Prior compensation, B pays cost } \chi_B} \tag{33}$$

where

$$\tilde{\pi}^{\text{B}}_{at|z_{bt}}(\text{w}) = E_{at|z_{bt}}[A_t]\text{w}^{\alpha} - \kappa\text{w} + \lambda_\tau \kappa\text{w}$$
$$\tilde{V}_{at|z_{bt}}(\text{w}) = E_{at|z_{bt}}\left[(1 - f_{\tau_t})[V(\mu_{at+1}, \tau + 1, \text{w})] + f_{\tau_t} V^R\right]$$

are the Board's flow utility (from company profits and bias) and continuation values upon receiving z_{bt}. The Board's proposed compensation policy $\tilde{\omega}(\mu_{at}, \sigma_{at}, w_{t-1})$ will in general differ from the baseline.[25]

I solve this version of the model using the estimated parameters and compare relevant outcomes to the baseline (using the same sequence of shocks in both simulations). Table 8 displays the results: the binding SOP failure rate and the average percentage changes in the CEO wage level and firm value, and these quantities are estimated and displayed for subsamples defined by observable characteristics. The SOP failure rate falls by a small amount, from 6.4% in the baseline to 6% in the counterfactual. Interestingly, however, total pay levels increase by about 5% and firm value decreases by about 1.5%

[25] Note that the adjustment cost is not present in failure by definition (compensation does not change from $t - 1$).

Figure 6. Impact of binding Say-on-Pay conditional on Shareholder belief precision

This figure illustrates the Shareholder's trade-off in the binding vote counterfactual between disciplining the Board and allowing the wage to fully reflect information about the CEO's ability. Using the simulated counterfactual data, each panel shows a binned scatterplot, where the main independent variable is the log change in the Shareholder's (prior) belief precision relative to CEO tenure $\tau = 1$. CEO fixed effects and a control for current (prior) beliefs about CEO ability are included in each regression. Because the binding vote is only relevant for tenure beyond the first year ($\tau > 0$), I drop $\tau = 0$ observations from the analysis. In Panel A, the dependent variable is the SOP failure rate; in Panel B, the counterfactual change in the wage level; and in Panel C, the counterfactual change in firm value.

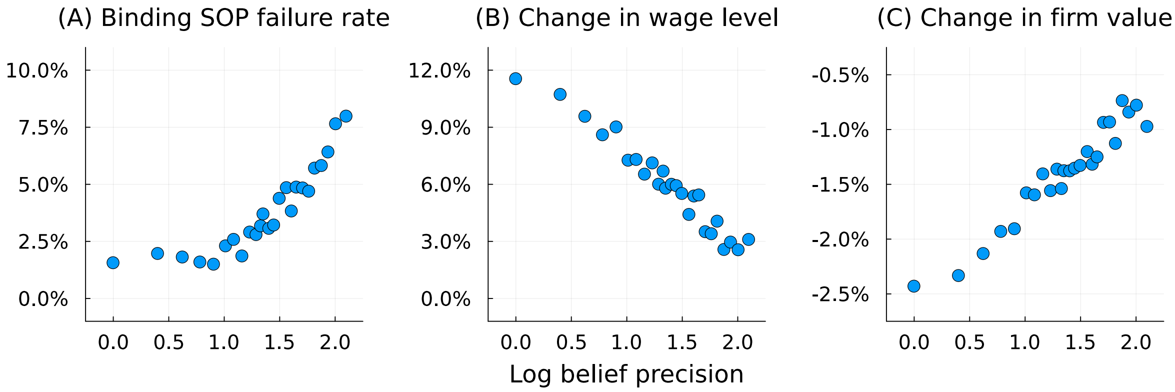

on average. The table shows that the binding vote reduces the effective disciplining power of the vote, even though it gives the Shareholder explicit control over compensation policy.

This arises because learning is valuable: the CEO receiving compensation that fully incorporates information often proves more beneficial than having direct control over compensation policy. Early in CEO tenure (when learning is most beneficial), the SOP failure rate would decrease to just 0.93%, while pay levels would increase by 10.3% (see Column 4 of Table 8). The Shareholder internalizes the benefit of an ability-appropriate wage early in tenure and optimally chooses a low failure rate. The Board anticipates and raises the proposed wage, thus lowering the power of SOP in mitigating board capture.

To illustrate further, Figure 6 displays how key outcomes change in the binding counterfactual as a function of Shareholder precision of beliefs about CEO ability. Using the counterfactual data, each panel shows a binned scatterplot in which the independent variable is the log precision of Shareholder beliefs about CEO ability (which strongly relates to the value of new information about CEO skill).

In Panel A, the binding failure rate is lower when beliefs are less precise, even after controlling for the CEO fixed effect (ability a), and current beliefs about ability (prior belief μ_{at}). When learning about CEO skill is most valuable (early in tenure), the Shareholder lowers the failure rate because the failing wage does not reflect current information about CEO ability. In turn, the Board optimally responds by raising the CEO's wage (by about 12%, relative to the baseline). As such, a binding SOP would lessen the effective disciplining power of the vote. The negative impact on firm value is shown in Panel C, where the counterfactual change in value tends toward zero as shareholder belief precision increases, though it remains negative across the entire range of belief precision.

My analysis broadly connects to Levit and Malenko (2011), who, from the standpoint of strategic voting, show how a binding vote may be less efficient at conveying shareholder views to management. It also connects to Levit (2020) who shows that the power of costly intervention may hinder communication between the shareholder and decision-maker as the decision-maker anticipates the intervention when setting the policy. I show that binding voting may lower realized disciplining power, if a failed vote binds the corporate policy to a state that has not aggregated all available information, precisely because the decision-maker internalizes that the Shareholder may value information over intervention.

7. Conclusion

CEO influence on the Board of Directors induces an agency problem in compensation policy. Say-on-Pay, *the* prominent shareholder voice mechanism in corporate governance, allows shareholders to discipline the Board and influence compensation policy. However, because of their non-binding nature and low failure rate, the impact of SOP votes is unclear.

This paper establishes this impact and clarifies the channel through which it occurs. Through my estimated structural model, I show that SOP resembles a costly punishment mechanism: shareholders can punish the Board, but doing so is costly. For the Board and shareholder, these costs are equivalent to 3.14% and 0.90% of each party's equity stake in the firm. The shareholder cost of giving punishment helps explain the low failure rate of SOP, but does not mitigate its value creation. I show that providing shareholders with the vote reduces wages by 6.6% and increases firm value by 2.4% on average.

I study the efficacy of non-binding *vs.* binding voting by constructing a counterfactual binding SOP. Vote failure binds CEO pay to its prior level, and though it entails greater discipline, failure implies that realized compensation policy may not reflect current information about CEO ability. I show that making SOP binding would actually diminish its effective disciplining power, especially when learning about CEO skill is valuable: shareholders value a compensation package more reflective of current information. These results have implications for the literature on shareholder democracy more generally.

My results broaden our understanding shareholder voting, suggesting that high support rates in shareholder votes may not imply ineffective monitoring: the threat of a failed vote may have a large impact, even if this impact is not directly observable. If corporate decision-makers internalize a cost to shareholder dissent on environmental or social shareholder proposals, then a mandated vote similar to SOP may positively impact shareholder welfare, even if the observed pass rate were high.

References

Aggarwal, R., Dahiya, S., and Prabhala, N. R. (2019). The power of shareholder votes: Evidence from uncontested director elections. *Journal of Financial Economics*, 133(1):134–153.

Aggarwal, R., Dahiya, S., and Yilmaz, U. (2023). Why do investors vote against corporate directors?

Allaire, Y. and Dauphin, F. (2016). Should say on pay votes be binding? https://clsbluesky.law.columbia.edu/2016/09/13/should-say-on-pay-votes-be-binding/ [Accessed: 2024-09-26].

Ambrus, A. and Greiner, B. (2012). Imperfect public monitoring with costly punishment: An experimental study. *American Economic Review*, 102(7):3317–3332.

Appel, I. R., Gormley, T. A., and Keim, D. B. (2016). Passive investors, not passive owners. *Journal of Financial Economics*, 121(1):111–141.

Armstrong, C. S., Gow, I. D., and Larcker, D. F. (2013). The efficacy of shareholder voting: Evidence from equity compensation plans. *Journal of Accounting Research*, 51(5):909–950.

Babenko, I., Choi, G., and Sen, R. (2019). Management (of) proposals.

Bach, L. and Metzger, D. (2019). How close are close shareholder votes? *The Review of Financial Studies*, 32(8):3183–3214.

Bebchuk, L. A. and Fried, J. M. (2003). Executive compensation as an agency problem. *Journal of economic perspectives*, 17(3):71–92.

Bertrand, M. and Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. *The Quarterly Journal of Economics*, 116(3):901–932.

Blonien, P., Crane, A. D., Crotty, K., and De Angelis, D. (2024). Proxy advice and errors in shareholder voting.

Brav, A., Jiang, W., Partnoy, F., and Thomas, R. (2008). Hedge fund activism, corporate governance, and firm performance. *The Journal of Finance*, 63(4):1729–1775.

Cai, J. and Walkling, R. A. (2011). Shareholders' say on pay: Does it create value? *Journal of Financial and Quantitative Analysis*, 46(2):299–339.

Cattaneo, M. D., Jansson, M., and Ma, X. (2018). Manipulation testing based on density discontinuity. *The Stata Journal*, 18(1):234–261.

Chaigneau, P., Edmans, A., and Gottlieb, D. (2022). A theory of fair ceo pay. *Working Paper*.

Coles, J. L., Daniel, N. D., and Naveen, L. (2014). Co-opted boards. *The Review of Financial Studies*, 27(6):1751–1796.

Correa, R. and Lel, U. (2016). Say on pay laws, executive compensation, pay slice, and firm valuation around the world. *Journal of Financial Economics*, 122(3):500–520.

Cuñat, V., Gine, M., and Guadalupe, M. (2012). The vote is cast: The effect of corporate governance on shareholder value. *The Journal of Finance*, 67(5):1943–1977.

Cuñat, V., Giné, M., and Guadalupe, M. (2016). Say pays! Shareholder voice and firm performance. *Review of Finance*, 20(5):1799–1834.

Dewatripont, M., Jewitt, I., and Tirole, J. (1999). The economics of career concerns, part I: Comparing information structures. *The Review of Economic Studies*, 66(1):183–198.

Dey, A., Starkweather, A., and White, J. T. (2024). Proxy advisory firms and corporate shareholder engagement. *The Review of Financial Studies*, 37(12):3877–3931.

Edmans, A., Gosling, T., and Jenter, D. (2023). CEO compensation: Evidence from the field. *Journal of Financial Economics*, 150(3):103718.

Federal Council, Switzerland (2013). Ordinance against excessive remuneration in listed companies limited by shares comes into force on 1 january 2014. https://www.admin.ch/gov/en/start/documentation/media-releases.msg-id-51029.html [Accessed 2024-11-01].

Ferri, F. (2015). Say on pay. In *Research Handbook on Shareholder Power*, pages 319–332. Edward Elgar Publishing.

Ferri, F. and Maber, D. A. (2013). Say on pay votes and CEO compensation: Evidence from the UK. *Review of Finance*, 17(2):527–563.

Fisch, J., Palia, D., and Solomon, S. D. (2018). Is say on pay all about pay: The impact of firm performance. *Harvard Business Law Review*, 8:101.

Flammer, C. (2015). Does corporate social responsibility lead to superior financial performance? A regression discontinuity approach. *Management science*, 61(11):2549–2568.

Fos, V. (2017). The disciplinary effects of proxy contests. *Management Science*, 63(3):655–671.

Fos, V. and Tsoutsoura, M. (2014). Shareholder democracy in play: Career consequences of proxy contests. *Journal of Financial Economics*, 114(2):316–340.

Gantchev, N., Gredil, O. R., and Jotikasthira, C. (2018). Governance under the gun: Spillover effects of hedge fund activism. *Review of Finance*, 23(6):1031–1068.

Gillan, S. L. and Starks, L. T. (2007). The evolution of shareholder activism in the United States. *Journal of Applied Corporate Finance*, 19(1).

Graham, J. R., Kim, H., and Leary, M. (2020). CEO-board dynamics. *Journal of Financial Economics*, 137(3):612–636.

Hartzell, J. C. and Starks, L. T. (2003). Institutional investors and executive compensation. *The Journal of Finance*, 58(6):2351–2374.

Hauder, E. A. (2019). Bouncing back from a low say-on-pay vote. https://corpgov.law.harvard.edu/2018/11/05/bouncing-back-from-a-low-say-on-pay-vote/ [Accessed: 2023-02-18].

Heath, D., Macciocchi, D., Michaely, R., and Ringgenberg, M. C. (2022). Do index funds monitor? *The Review of Financial Studies*, 35(1):91–131.

Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*, volume 25. Harvard University Press.

Holderness, C. G. (2018). Equity issuances and agency costs: The telling story of shareholder approval around the world. *Journal of Financial Economics*, 129(3):415–439.

Holland, K., Lim, C., and Yi, I. (2023). Shareholder meeting uncertainty: Evidence from the options market.

ISS (2022). United States proxy voting guidelines. https://www.issgovernance.com/file/policy/active/americas/US-Voting-Guidelines.pdf [Accessed: 2023-02-25].

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review*, 76(2):323–329.

Johnson, T. L. and Swem, N. (2021). Reputation and investor activism: A structural approach. *Journal of Financial Economics*, 139(1):29–56.

Kakhbod, A., Loginova, U., Malenko, A., and Malenko, N. (2023). Advising the management: A theory of shareholder engagement. *The Review of Financial Studies*, 36(4):1319–1363.

Kaplan, S. N. (2013). CEO pay and corporate governance in the US: Perceptions, facts, and challenges. *Journal of Applied Corporate Finance*, 25(2):8–25.

Levit, D. (2019). Soft shareholder activism. *The Review of Financial Studies*, 32(7):2775–2808.

Levit, D. (2020). Words speak louder without actions. *The Journal of Finance*, 75(1):91–131.

Levit, D. and Malenko, N. (2011). Nonbinding voting for shareholder proposals. *The Journal of Finance*, 66(5):1579–1614.

Listokin, Y. (2008). Management always wins the close ones. *American Law and Economics Review*, 10(2):159–184.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, pages 995–1026.

Mehran, H. (1995). Executive compensation structure, ownership, and firm performance. *Journal of Financial Economics*, 38(2):163–184.

Netflix, Inc. (2022). Netflix 2022 proxy statement. https://www.sec.gov/Archives/edgar/data/1065280/000119312522115102/d202881ddef14a.htm#toc202881_21 [Accessed: 2023-02-23].

Page, T. B. (2018). CEO attributes, compensation, and firm value: Evidence from a structural estimation. *Journal of Financial Economics*, 128(2):378–401.

Pinnington, J. (2023). Are passive investors biased voters?

Shleifer, A. and Vishny, R. W. (1997). A survey of corporate governance. *The Journal of Finance*, 52(2):737–783.

Silveira, B. S. (2017). Bargaining with asymmetric information: An empirical study of plea negotiations. *Econometrica*, 85(2):419–452.

Taylor, L. A. (2010). Why are CEOs rarely fired? Evidence from structural estimation. *The Journal of Finance*, 65(6):2051–2087.

Taylor, L. A. (2013). CEO wage dynamics: Estimates from a learning model. *Journal of Financial Economics*, 108(1):79–98.

Zhu, C. (2021). The preventive effect of hedge fund activism: Investment, CEO compensation and payout policies. *International Journal of Managerial Finance*, 17(3):401–415.

Table 1. Summary statistics

This table displays descriptive statistics of the variables used in the empirical analysis and estimating the model. The sample is based upon a merge of Compustat, Execucomp and ISS for the years 2011-2020. All dollar variables are deflated to 2018 dollars. There are a total of 2,112 CEO spells in the dataset.

	N	Mean	Std dev	25%	50%	75%
Say-on-Pay						
Vote against SOP (%)	10,170	0.087	0.121	0.024	0.043	0.084
SOP failure: $\mathbb{1}$[More than 30% against]	10,170	0.070				
$\mathbb{1}$[More than 20% against]	10,170	0.117				
$\mathbb{1}$[More than 50% against]	10,170	0.020				
CEO						
Total compensation (TDC1, $m)	10,170	7.268	7.018	2.961	5.533	9.537
Current compensation (share, %)	10,170	0.272	0.215	0.134	0.200	0.325
Change in total pay (%)	7,903	0.040	0.460	-0.099	0.046	0.198
CEO tenure (years)	10,170	7.062	5.603	3	6	10
Length of CEO tenure (years)	2,112	9.321	5.878	5	8	13
Firm						
Assets ($b)	10,170	28.754	154.601	1.202	3.830	12.710
Revenues ($b)	10,170	9.144	25.058	0.774	2.194	6.989
Market capitalizaiton ($b)	10,170	15.571	55.359	1.177	3.137	10.838
Return on assets (%)	10,170	0.120	0.093	0.067	0.114	0.165
Profit margin (%)	10,170	0.197	0.165	0.097	0.168	0.283

Table 2. Say-on-Pay results, CEO pay and company performance

This table explores the relation between SOP outcomes, total CEO pay and company performance, in support of Fact 1. Panel A estimates the relation between SOP outcomes and the level of CEO pay. The dependent variable in columns 1-3 is an indicator for SOP failure (at least 30% of shareholders voting against the SOP), and in columns 4-6 it is the percentage of shareholders that vote against. Log total pay is the natural logarithm of the CEO's total expected pay (TDC1). SOP fail is an indicator if a SOP vote fails, i.e. the % voting against is above 30%. % vote against in SOP is the proportion of shareholders voting to fail the SOP. Return on assets is standardized to mean zero, unit variance. Standard errors are displayed below coefficients and clustered at the firm × CEO level. ***, **, * denote significance at 1%, 5%, 10%.

	(1)	(2)	(3)	(4)	(5)	(6)
	SOP fail {0,1}			% vote against in SOP		
Log total pay	0.071***	0.077***	0.078***	0.035***	0.039***	0.039***
	(0.008)	(0.009)	(0.009)	(0.004)	(0.004)	(0.004)
Return on assets		-0.024***	-0.025***		-0.016***	-0.016***
		(0.007)	(0.007)		(0.003)	(0.003)
Log assets		0.006	-0.002		0.001	-0.003
		(0.014)	(0.014)		(0.006)	(0.006)
Lagged log total pay			0.035***			0.016***
			(0.008)			(0.004)
Observations	10,170	10,170	10,170	10,170	10,170	10,170
R-squared	0.398	0.404	0.406	0.491	0.498	0.500
Firm × CEO FE	Yes	Yes	Yes	Yes	Yes	Yes
CEO tenure FE		Yes	Yes		Yes	Yes
Year FE		Yes	Yes		Yes	Yes

Table 3. Evidence of costs to directors from Say-on-Pay failure

This table displays correlations between costly outcomes for directors and SOP failure. I focus on directors that serve on the compensation committee during the year of an SOP vote. Director turnover (columns 1-2) occurs when a director is no longer on the Board the year after an SOP vote, conditional on the next year not being the year the director's term ends. Compensation committee turnover (columns 3-4) occurs when a compensation committee director is no longer on the compensation committee the year after an SOP vote, conditional on the director remaining on the Board. A reduction in outside Board positions (columns 5-6) occurs when the number of outside boards a director sits on decreases in the year after an SOP vote, conditional on the director sitting on at least one outside Board and the director remaining on their current Board the next year. SOP fail is an indicator if a SOP vote fails, i.e. the % voting against is above 30%. % vote no in SOP is the proportion of shareholders voting to fail the SOP. The firm's return on assets is included as a covariate in columns 1-4, the mean outside firms' return on assets is included as a covariate in columns 5-6. All continuous covariates are standardized to mean zero, unit variance. Standard errors are displayed below coefficients and clustered at the director level. ***, **, * denote significance at 1%, 5%, 10%.

	(1)	(2)	(3)	(4)	(5)	(6)
	Turnover		Off comp. committee		Reduction outside seats	
SOP fail {0,1}	2.226***	1.619***	1.387***	1.353***	1.922**	1.722**
	(0.539)	(0.512)	(0.449)	(0.459)	(0.854)	(0.860)
Log director age		2.398***		-0.230		1.935***
		(0.184)		(0.152)		(0.290)
Log board size		0.340**		0.334**		0.203
		(0.160)		(0.140)		(0.257)
Log director tenure		1.361***		0.192		-0.187
		(0.159)		(0.145)		(0.274)
Log CEO tenure		-1.206***		-0.800***		
		(0.163)		(0.140)		
Log total CEO pay		0.545***		0.065		
		(0.150)		(0.156)		
Constant	10.147***		5.481***		9.488***	
	(0.173)		(0.145)		(0.293)	
Observations	33,213	33,213	29,752	29,752	13,469	13,469
R-squared	0.001	0.186	0.001	0.009	0.000	0.006
Year FE		Yes		Yes		Yes
Industry FE		Yes		Yes		Yes

Table 4. Model fit

This table displays the estimated model's fit. Following the identification strategy in Section 4.1, I estimate the model from Section 3 using the simulated method of moments, which is described in detail in Appendix IA2. I display data and model moments (under the estimated optimal parameter vector vector), the difference between the moments, and a t-test of the difference between the observed and simulated data moments.

	Description	Notation	Data: Observed	Data: Simulated	Difference	t-stat
(1)	SOP failure rate	s_0	0.068	0.064	0.004	0.391
(2)	SOP fail wage sensitivity	s_1	0.048	0.054	-0.006	-0.371
(3)	SOP fail output sensitivity	s_2	-0.022	-0.027	0.005	0.177
(4)	SOP failure residual variance	$Var(\epsilon^s)$	0.038	0.046	-0.008	-3.775
(5)	Average log wage SOP pass	b_0	1.604	1.599	0.005	4.052
(6)	Change in log wage SOP fail	b_1	0.221	0.200	0.021	0.564
(7)	Wage growth over tenure	b_2	0.049	0.008	0.040	8.611
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.114	0.115	-0.002	-0.056
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	0.596	0.596	0.000	0.014
(10)	Average log output	y_0	7.717	7.775	-0.058	-1.376
(11)	Elasticity of output to wage	y_1	1.191	1.193	-0.002	-0.537
(12)	Output residual variance	$Var(\epsilon^y)$	2.446	2.403	0.042	10.294
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	2.253	2.255	-0.002	-1.078
(14)	Average profit margin	Π / y	0.213	0.282	-0.069	-3.985

Table 5. Estimated parameters

This table displays estimates of the parameters that drive the model in Section 3. The panel also displays the magnitudes of the Board and Shareholder costs to SOP failure as a percentage of model-implied average firm value. I compute average firm value in closed form as a function of model parameters, and use the delta method to calculate standard errors (see Appendix IA1.5 for derivation).

Description	Notation	Value
SOP failure costs		
Board SOP failure cost (% average firm value)	χ_B / V_0	3.14% (0.112%)
Shareholder SOP failure cost (% average firm value)	χ_S / V_0	0.90% (0.035%)
Estimated parameters		
Board SOP failure cost	χ_B	2.247 (0.0655)
Shareholder SOP failure cost	χ_S	0.641 (0.0292)
CEO board capture (constant)	λ_0	0.131 (0.0050)
CEO board capture (growth)	λ_1	0.066 (0.0085)
Prior std dev of CEO ability	σ_0	1.105 (0.0083)
Output—CEO wage elasticity	α	0.286 (0.0039)
Std dev of productivity shock	σ_y	1.465 (0.0048)
Scaling factor (output)	$\log \eta$	7.321 (0.0434)
Scaling factor (cost)	κ	0.329 (0.0067)
Std dev of Board signal	σ_{z_b}	0.818 (0.0131)
Std dev of Shareholder signal	σ_{z_s}	0.684 (0.0067)
CEO wage adjustment cost	c_w	4.857 (0.0356)

Table 6. The impact of Say-on-Pay on compensation policy and firm value

This table analyzes the impact of SOP on compensation policy and firm value. I simulate a counterfactual (using the same sequence of shocks as in the baseline estimation) in which the Board cost to SOP failure (χ_B) is set to zero while holding other parameters constant, effectively analyzing compensation policy as if SOP did not exist. The first row displays the average counterfactual percentage change in the wage level; the second displays the change in objective firm value from (22). The first column displays the average percentage changes for the entire simulated sample. The remaining columns split the sample based on cross-sectional characteristics. In columns two and three, "Low Ability" means that the CEO's type a is less than zero (i.e., below the average); "High" ability means that the CEO's type a is greater than zero. In columns four and five, "Early" refers to first five years of tenure ("Late" refers to years of tenure beyond five). In columns six and seven, "Fail" refers to observations in which the SOP vote would fail in the baseline; "Pass" refers the opposite.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Ability		**Tenure**		**SOP vote**	
	Full sample	Low	High	Early	Late	Fail	Pass
% change in							
Wage level	-6.62%	-8.45%	-4.83%	-3.57%	-10.26%	-13.83%	-6.06%
Firm value	+2.35%	+3.44%	+1.24%	+0.64%	+4.73%	+9.83%	+1.86%

Table 7. Subsample heterogeneity

This table presents estimation outcomes for subsamples split by various characteristics, using the same routine as the main sample (see Table 5). For "Board co-option" is the percentage of directors appointed during the CEO's tenure, adjusted for tenure effects (Coles et al., 2014). "Top 5 inst. ownership" is measured as the percentage of equity held by the five largest investors (Hartzell and Starks, 2003). Board co-option compares the bottom and top quintiles, whereas the top 5 inst. ownership compares below- and above-median. SOP failure costs and parameters from each split and t-statistics testing parameter equality are shown.

		(1)	(2)	(3)	(4)	(5)	(6)
		Board co-option			Top 5 inst. ownership		
	Notation	**Low**	**High**	***t*-stat**	**Low**	**High**	***t*-stat**
SOP failure costs							
(1)	χ_B / V_0	4.28%	1.33%	21.32%	3.62%	4.48%	-2.86%
(2)	χ_S / V_0	1.49%	1.58%	-0.05%	1.00%	2.66%	-26.24%
Estimated parameters							
(3)	χ_B	2.233	1.435	26.247	2.327	3.107	-12.057
(4)	χ_S	0.776	1.712	-0.971	0.645	1.848	-19.470
(5)	λ_0	0.126	0.099	4.416	0.069	0.185	-6.336
(6)	λ_1	0.082	0.105	-3.922	0.069	0.077	-0.561
(7)	σ_0	1.022	1.353	-18.148	0.997	1.238	-9.821
(8)	α	0.300	0.294	0.373	0.365	0.278	2.078
(9)	σ_y	1.358	1.460	-6.286	1.389	1.380	0.843
(10)	$\log \eta$	7.225	7.297	-0.358	7.958	6.997	8.395
(11)	κ	0.340	0.351	-1.126	0.357	0.377	-0.718
(12)	σ_{z_b}	0.646	0.441	21.735	0.584	0.424	17.844
(13)	σ_{z_s}	1.040	0.527	73.100	1.451	1.918	-34.362
(14)	c_w	5.395	6.002	-4.792	0.555	7.906	-161.220

Table 8. Counterfactual exercise: Binding Say-on-Pay

This table displays a counterfactual experiment where I re-solve a different version of the model, in which SOP votes are binding and failure implies that the CEO's pay is set to its previous value w_{t-1}. The first row displays the counterfactual SOP failure rate. Rows 2 and 3 display the counterfactual percentage change in wages and firm value. The first column displays for the full sample, columns 2-7 are the same cross-sectional splits as in Table 6. To compute the counterfactual, I re-solve the counterfactual model, applying the same sequence of shocks to each firm. I solve for optimal choices, and solve for the percentage change in each quantity at the observation level. Because the binding vote is only relevant for tenure beyond the first year ($\tau > 0$), I drop $\tau = 0$ observations from the analysis.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Ability		**Tenure**		**SOP vote**	
	Full sample	Low	High	Early	Late	Fail	Pass
SOP fail %							
Non-binding	6.43%	11.80%	3.21%	3.62%	10.38%	100.00%	0.00%
Binding	5.99%	7.41%	5.16%	0.93%	12.25%	46.85%	2.77%
% change in							
Wage level	+5.23%	+7.37%	+4.27%	+10.34%	+3.33%	+0.06%	+5.85%
Firm value	-1.52%	-1.70%	-1.65%	-1.93%	-1.86%	-0.25%	-1.61%

A. Appendix Figures and Tables

Figure A1. Example of monitoring costs arising from SOP failure — 2021 Netflix SOP

This figure presents anecdotal evidence of monitoring costs incurred by shareholders when the SOP fails. The 2021 Netflix SOP saw 49.4% of shares voting against the SOP. Under the 30% failure rule, this presents a clear SOP failure. Netflix directors then repeatedly engaged with large stockholders in the following year over the compensation policy. See Netflix, Inc. (2022).

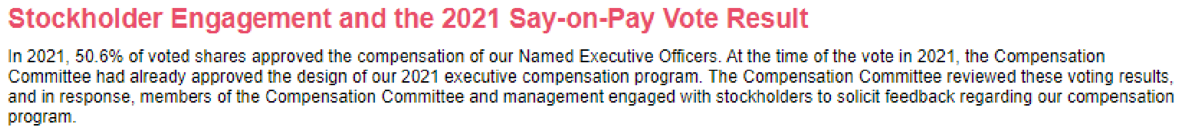
Stockholder Engagement and the 2021 Say-on-Pay Vote Result

In 2021, 50.6% of voted shares approved the compensation of our Named Executive Officers. At the time of the vote in 2021, the Compensation Committee had already approved the design of our 2021 executive compensation program. The Compensation Committee reviewed these voting results, and in response, members of the Compensation Committee and management engaged with stockholders to solicit feedback regarding our compensation program.

Figure A2. Bunching at SOP vote thresholds of 30% and 50%

This figure displays the result of testing for density manipulation of SOP votes at the failure thresholds of $k = \{30\%, 50\%\}$, following the methodology described in Cattaneo et al. (2018). I focus on SOP votes falling within 10 pp of each failure threshold and test for density manipulation at the failure threshold. The blue and orange bars display observed frequencies of Δ^{data} in 0.5 percentage point bins and the blue and orange lines display the estimated density.

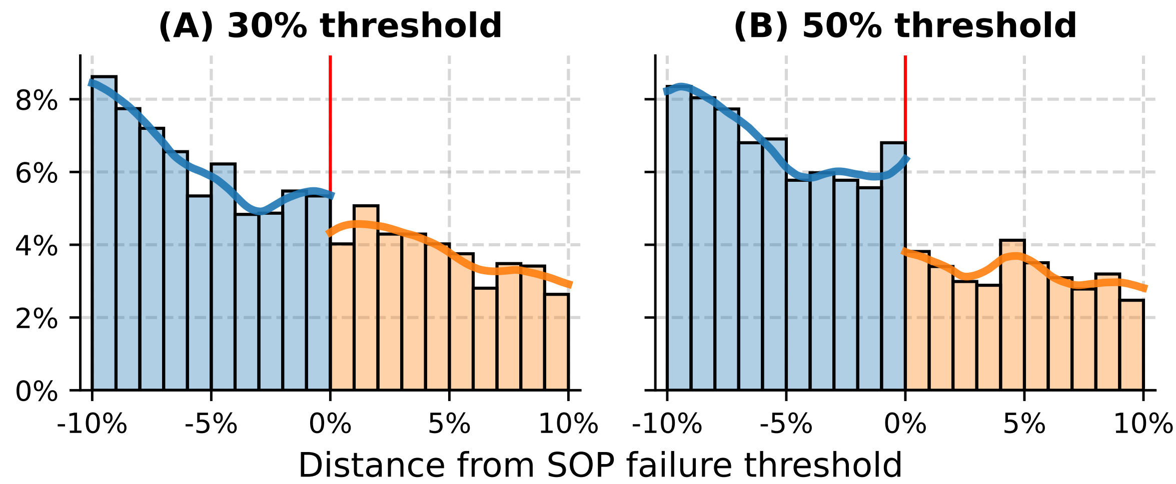

Figure A3. Board and Shareholder volatility of beliefs over CEO tenure

This figure displays the volatility Board and shareholder beliefs about CEO ability as a function of CEO tenure τ at the time at which each agent plays their strategy, using the parameter estimates from Table 5. Each volatility of beliefs is at the time that each party plays their strategy: the Board's belief volatility is $\sigma_{bt|z_b}$, as in (19), and the Shareholder's belief volatility is σ_{at} (i.e., the prior).

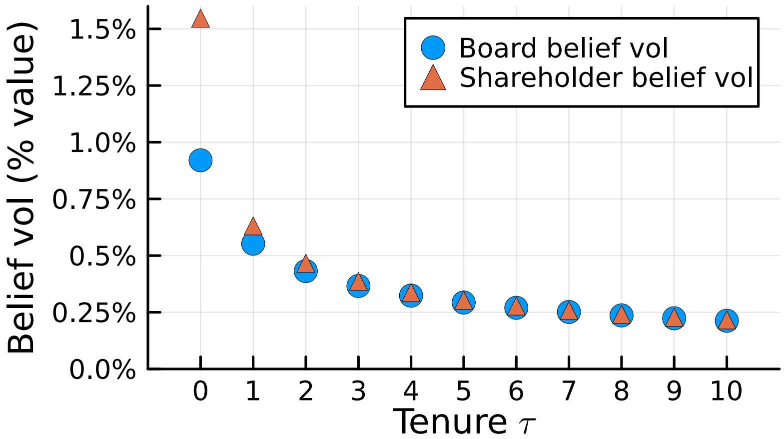

Table A1. Changes in CEO compensation following SOP disapproval

This table explores the relation between changes in CEO compensation (from t to $t+1$) and SOP results (from t). The dependent variable is the log change in CEO compensation from t to $t+1$. SOP fail is an indicator if a SOP vote fails, i.e. the % voting against is above 30%. % vote no in SOP is the proportion of shareholders voting to fail the SOP. Return on assets is standardized to mean zero, unit variance. Standard errors are displayed below coefficients and clustered at the firm × CEO level. ***, **, * denote significance at 1%, 5%, 10%.

	(1)	(2)	(3)	(4)	(5)	(6)
	Change in total pay (t to $t+1$)					
SOP fail {0,1}	-0.086***	-0.081***	-0.081***			
	(0.030)	(0.029)	(0.029)			
% vote against in SOP				-0.234***	-0.177***	-0.177***
				(0.071)	(0.067)	(0.067)
Log total pay	-0.951***	-1.049***	-1.049***	-0.948***	-1.047***	-1.047***
	(0.025)	(0.024)	(0.024)	(0.025)	(0.025)	(0.025)
Return on assets		0.035**	0.035**		0.034**	0.034**
		(0.016)	(0.016)		(0.016)	(0.016)
Log assets		0.125***	0.125***		0.124***	0.124***
		(0.029)	(0.029)		(0.029)	(0.029)
Observations	7,502	7,502	7,502	7,502	7,502	7,502
R-squared	0.540	0.580	0.580	0.540	0.580	0.580
Firm × CEO FE	Yes	Yes	Yes	Yes	Yes	Yes
CEO tenure FE		Yes	Yes		Yes	Yes
Year FE		Yes	Yes		Yes	Yes

Table A2. Board co-option and the level of CEO pay

This table displays how board co-option (Coles et al., 2014), an important empirical measure of *board capture* influences the level of CEO pay and modulates the effect of SOP results on changes in CEO compensation, in support of Fact 2. Panel A presents correlations between the level of CEO pay and the degree of board co-option. Board co-option and Return on assets are standardized to mean zero, unit variance. Standard errors are displayed below coefficients and clustered at the firm × CEO level. ***, **, * denote significance at 1%, 5%, 10%.

	(1)	(2)	(3)	(4)
	Log total pay			
Board co-option	0.078***	0.089***	0.086***	0.061***
	(0.028)	(0.020)	(0.019)	(0.010)
Return on assets		0.235***	0.236***	0.103***
		(0.018)	(0.018)	(0.009)
Log assets		0.325***	0.343***	0.389***
		(0.008)	(0.011)	(0.006)
Log board size			-0.232***	0.035
			(0.086)	(0.044)
Observations	6,347	6,347	6,347	6,346
R-squared	0.014	0.441	0.444	0.616
CEO tenure FE	Yes	Yes	Yes	Yes
Industry FE				Yes

Internet Appendix

IA1. Model Appendix

IA1.1. Microfoundation of Representative Shareholder Assumption

This section microfounds the assumption of a representative shareholder, under the assumption that a diffuse shareholder base receive a correlated private signal of CEO ability (see Section 3.2). When all atomistic shareholders receive a correlated signal and vote with the same threshold strategy, it is informationally equivalent to focus on a representative shareholder votes with the same threshold strategy under the aggregated signal. This assumption of correlated private signals also embeds proxy advisors into the model: if a proxy advisor gives a negative recommendation, this is like a strong, correlated negative signal of CEO ability.

Proposition 1. *The expected proportion of shareholders voting against the SOP is informationally equivalent to*

$$CDF_s^U(k_s(w))$$

where $CDF_s^U(\cdot)$ is the CDF of the distribution of the random variable which determines the outcome of the SOP vote and $k_s(w) = s \times w$, where w is the CEO's wage and s is the shareholder's choice variable.

Proof. The proof largely follows arguments in Pinnington (2023). In the model, there is a continuum of N_S shareholders, whom each draw a signal z_{si} that is private knowledge, but correlated across shareholders,

$$z_{si} = \overline{z} + \varepsilon_{si}, \quad \varepsilon_{si} \sim N\left(0, \sigma_{si}^2\right).$$

z_i is conditionally normal and independent across shareholders given the common, latent signal $\overline{z}$, distributed according to

$$\overline{z} = a + \varepsilon_{\overline{z}}, \quad \varepsilon_{\overline{z}} \sim N\left(0, \sigma_{\overline{z}}^2\right).$$

The standard voting model with incomplete information assumes that signals are completely private, i.e $z_{it} = a + \varepsilon_{zit}$. With proxy voting, signals are more likely to be correlated. For example, $\overline{z}$ could reflect proxy advisors' recommendations. Note, however, that $\overline{z}$ is not a public signal. Rather, each shareholder shares the same belief about $\overline{z}$. So, it is as if shareholders each receive the proxy advisor's signal with some "noise," which could reflect, e.g. idiosyncratic trust in the proxy advisor across shareholders.

Shareholders play a symmetric cutoff strategy, voting against the proposal if and only if they draw a signal below their cutoff value

$$\mathbb{1}[\text{SOP fail}_i] \iff z_i \leq k_s^i(w).$$

Note — I have abstracted away from the effect of the output shock on the vote, and adjudging failure using lognormals. Given that the output shock is common knowledge, it will affect all shareholers voting in the same way, so does not impact the proof; the conversion to lognormal is a technical assumption that again affects all shareholders equivalently.

Given $\bar{z}$, the probability that a single shareholder votes against is

$$\Pr(\text{SOP fail}_i \mid \bar{z}) = \Phi\left(\frac{k_s^i(w) - \bar{z}}{\sigma_{si}}\right)$$

and the probability I observe N out of N_S shareholders voting against is

$$Pr(N \mid \bar{z}) = C_N^{N_S}\left[\Phi\left(\frac{k_s^i(w) - \bar{z}}{\sigma_{si}}\right)\right]^N \left[1 - \Phi\left(\frac{k_s^i(w) - \bar{z}}{\sigma_{si}}\right)\right]^{N_S - N}.$$

Fixing the unknown type a, I can find the probability of observing N out of N_S against votes,

$$Pr(N \mid a) = \int f(\bar{z} \mid a) Pr(N \mid \bar{z})\, d\bar{z}.$$

Let p be the proportion of shareholders voting against: $p = N/N_S$. Since p is Binomial, as $N_S \longrightarrow \infty$, the distribution of p becomes increasingly peaked around its mean. Since its mean is the probability any individual shareholder votes against the proposal, the likelihood of observing p vanishes in the limit when $Pr(SOP_i = 1 \mid \bar{z})$ is anything other than p. Given that $\bar{z}$ completely determines $Pr(SOP_i = 1 \mid \bar{z})$, there is a bijection between $\bar{z}$ and p

$$\bar{z}(p) = k_s(w) - \sigma_{si}\Phi^{-1}(p)$$

Using this peakedness, the limit of the density of observing p as $N_S \longrightarrow \infty$ is

$$f(p) = \int f(a) f(\bar{z}(p) \mid a) \bar{z}'(p)\, da$$

The likelihood of observing p is driven by the likelihood of observing $\bar{z}(p)$, scaled by a change-of-variable term $\bar{z}'(p)$. Since $\bar{z}$ is conditionally normal around the type a, I integrate over all types a and then take the likelihood of observing $\bar{z}(p)$ given the type a. I am more interested in $f(a \mid p)$ — the density of a conditional on observing p,

$$\lim_{N_S \to \infty} f(a \mid p) = \frac{\lim_{N_S \to \infty} f(a) f(p \mid a)}{\int \lim_{N_S \to \infty} f(a) f(p \mid a)\, da}.$$

The intuition is that the posterior likelihood of a is proportional to two components: the prior $f(a)$; and the likelihood that the latent signal $\bar{z}$, given a, is equal to $\bar{z}(p)$, which in the limit is the only $\bar{z}$ for which I would see p. This is a scaled product of Gaussians, so the posterior is also normal,

$$a \mid p \sim N(\mu_{ap}, \sigma_{ap}^2),$$

where

$$\mu_{ap} = \frac{\sigma_{\bar{z}}^2}{\sigma_a^2 + \sigma_{\bar{z}}^2} \mu_a + \frac{\sigma_a^2}{\sigma_a^2 + \sigma_{\bar{z}}^2} \bar{z}(p), \qquad \sigma_{ap}^2 = \frac{\sigma_a^2 \sigma_{\bar{z}}^2}{\sigma_a^2 + \sigma_{\bar{z}}^2}.$$

Thus, observing p is informationally equivalent to observing a signal $z_s = \bar{z}(p) = k_s(w) - \sigma_z \Phi^{-1}(p)$, where $z_s = a + \varepsilon_s$, and $\varepsilon_s \sim N(0, \sigma_{z_s}^2)$. The proof arises because of the assumptions about the correlated signal and the continuum of shareholders. All shareholders play a symmetric cutoff strategy; in the limit, the exact proportion of shareholders that receive a signal below the cutoff must be equivalent to the probability that an informationally equivalent aggregate signal falls below the cutoff. Another way to think about this is to consider a representative shareholder that interacts with the Board, and aggregates the votes or signals of the shareholder base at the shareholder meeting. ■

IA1.2. Evolution of Board and Shareholder Beliefs

I first detail two Propositions, which define how beliefs update in the model. Then I define exactly how Board and shareholder beliefs change within each period.

IA1.2.1. Evolution of Beliefs Period to Period

Prop. 2 shows how beliefs change from t to $t+1$. Prop. 3 describes the distribution of next period beliefs given today's beliefs, which is used when the Board calculates their (expected) continuation value. Figure IA1 displays a detailed timeline that illustrated belief updating and the timing of strategies.

Proposition 2. *From period t to $t+1$, the variance of beliefs for both the Board and shareholders declines*

Figure IA1. Model timeline with exact time that strategies are played

This figure displays a more detailed model timeline, with mappings to the relevant assumptions, and when the Board and Shareholder play their strategies. See Figure 2 in the main text for the timeline as it maps to real-world outcomes. See Appendix IA1.3.

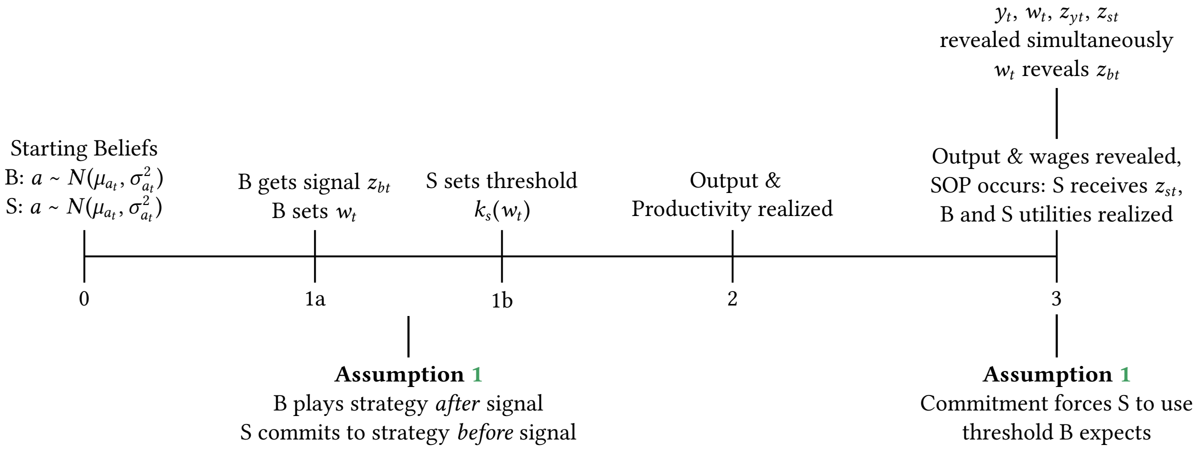

deterministically according to

$$\sigma_a^2(\tau+1) = \left[\sigma_a^{-2}(\tau) + \sigma_{z_b}^{-2} + \sigma_{z_s}^{-2} + \sigma_y^{-2}\right]^{-1} \quad \text{(IA1)}$$

where τ is the tenure of the CEO at year t. Equivalently, I can write the variance of beliefs about CEO ability entirely as a function of CEO tenure τ and model parameters

$$\sigma_a^2(\tau) = \sigma_0^2\left[1 + \tau\left(\kappa_{z_b}^{-1} + \kappa_{z_s}^{-1} + \kappa_y^{-1}\right)\right]^{-1} \quad \text{(IA2)}$$

where $\kappa_{z_b} = \sigma_{z_b}^2/\sigma_0^2$, $\kappa_{z_s} = \sigma_{z_s}^2/\sigma_0^2$ and $\kappa_y = \sigma_y^2/\sigma_0^2$

Similarly, from period t to $t+1$, the mean of beliefs for both the Board and shareholders evolves according to

$$\mu_{at+1} = \sigma_a^2(\tau+1)\left[\frac{\mu_{at}}{\sigma_a^2(\tau)} + \frac{z_{bt}}{\sigma_{z_b}^2} + \frac{z_{st}}{\sigma_{z_s}^2} + \frac{z_{yt}}{\sigma_y^2}\right] \quad \text{(IA3)}$$

Proof. The formulas are standard results in Bayesian learning (e.g., Pastor and Veronesi, 2009; Taylor, 2010).[1] The Board and shareholder reveal their signals each period. Thus, Board and shareholders share the same beliefs about the variance from period to period. ■

[1] See also the internet appendix for Taylor (2010).

Proposition 3. *The mean and variance of the mean of t + 1 CEO beliefs at t are*

$$E_t[\mu_{at+1}] = \mu_{at}$$
$$Var_t[\mu_{at+1}] = \sigma_a^2(\tau) - \sigma_a^2(\tau+1) \tag{IA4}$$

That is,

$$\mu_{at+1} \mid \mu_{at}, \tau \sim N\big(\mu_{at}, \sigma_a^2(\tau) - \sigma_a^2(\tau+1)\big)$$

Proof. I drop time subscripts for convenience, and use $\prime$ to denote next period. Via Prop. 2, the mean evolves as

$$\mu_{a\prime} = \sigma_{a\prime}^2\left[\frac{\mu_a}{\sigma_a^2} + \frac{z_b}{\sigma_{z_b}^2} + \frac{z_s}{\sigma_{z_s}^2} + \frac{z_y}{\sigma_y^2}\right]$$

where $z_y = \ln A = a + \varepsilon_y$ is the productivity signal. Let $p' = \sigma_{a'}^{-2}$, i.e the next period precision of beliefs. Let p_a, p_b, p_s, p_y be precisions $\sigma_a^{-2}, \sigma_{z_b}^{-2}, \sigma_{z_s}^{-2}, \sigma_y^{-2}$ respectively. Then define $\rho_{X\in\{a,b,s,y\}}$ be each precision divided by p', e.g. $\rho_a = \frac{p_a}{p'}$. I can write,

$$E[\mu_{a'} \mid \mu_a] = \big(\rho_a + \rho_b + \rho_s + \rho_y\big)\mu_a = \mu_a$$

which holds because the sum of precisions equals next period's precision. I can write $Var(\mu_{a'} \mid \mu_a)$ as

$$\begin{aligned}
Var(\mu_{a'} \mid \mu_a) &= E\Big[\big(\rho_a\mu_a + \rho_b z_b + \rho_s z_s + \rho_y z_y - E[\mu_{a'} \mid \mu_a]\big)^2 \mid \mu_a\Big] \\
&= E\Big[\big(\rho_a(\mu_a - \mu_a) + \rho_b(z_b - \mu_a) + \rho_s(z_s - \mu_a) + \rho_y\big(z_y - \mu_a\big)\big)^2 \mid \mu_a\Big] \\
&= E\Big[\big(\rho_b(z_b - \mu_a) + \rho_s(z_s - \mu_a) + \rho_y\big(z_y - \mu_a\big)\big)^2 \mid \mu_a\Big]
\end{aligned}$$

Note that $E\big[(z_b - \mu_a)^2 \mid \mu_a\big] = \sigma_a^2 + \sigma_{z_b}^2$, which similarly holds for subscript s and y. Hence, I can write $Var(\mu_{a'} \mid \mu_a)$ as

$$Var(\mu_{a'} \mid \mu_a) = \sigma_a^2\big(\rho_b + \rho_s + \rho_y\big)^2 + \rho_b^2\sigma_{z_b}^2 + \rho_s\sigma_{z_s}^2 + \rho_y\sigma_y^2$$

Note that $\rho_b^2\sigma_{z_b}^2 = \frac{\rho_b}{p'}$, similarly for s and y, and $1 = \rho_a + \rho_b + \rho_s + \rho_y$, so I can again write

$$Var(\mu_{a'} \mid \mu_a) = \sigma_a^2(1 - \rho_a)^2 + \frac{1 - \rho_a}{p'}$$

Lastly, I note that $\sigma_a^2(1-\rho_a) = \sigma_a^2 - \frac{\sigma_a^2 p_a}{p'} = \sigma_a^2 - \sigma_{a'}^2$, and

$$\begin{aligned} Var(\mu_{a'} \mid \mu_a) &= (\sigma_a^2 - \sigma_{a'}^2)(1-\rho_a) + \sigma_{a'}^2(1-\rho_a) \\ &= \sigma_a^2(1-\rho_a) \\ &= \sigma_a^2 - \sigma_{a'}^2 \end{aligned}$$

This expression of the conditional variance of the mean can be used for pair of normal prior + posterior beliefs. This quantity is useful when taking expectation of next period's continuation value ■

IA1.2.2. Differences in Board and Shareholder Beliefs Within Period

This sections how Board and shareholder beliefs evolve within each period. As the wage and vote reveal signals z_b and z_s, B and S share the same beliefs at the beginning of t. By Prop. 2, I have that $\sigma_{bt}^2 = \sigma_{st}^2 = \sigma_a^2(\tau_t)$ from (IA2). Let μ_{at} be the beliefs about the mean at the beginning of period t. The evolution of B and S beliefs within t is:

1. **Board beliefs after the compensation committee meeting**

   At the meeting, the Board receives signal z_{bt}, and Board beliefs update to

$$\mu_{bt|z_b} = \sigma_{bt|z_b}^2 \left(\frac{\mu_{at}}{\sigma_a^2(\tau_t)} + \frac{z_{bt}}{\sigma_{z_b}^2} \right) \tag{IA5}$$

$$\sigma_{bt|z_b}^2 = \frac{\sigma_a^2(\tau_t)\sigma_{z_b}^2}{\sigma_a^2(\tau_t) + \sigma_{z_b}^2} = \sigma_0^2 \left[1 + (\tau_t + 1)\kappa_{z_b}^{-1} + \tau_t(\kappa_{z_s}^{-1} + \kappa_y^{-1}) \right]^{-1} \tag{IA6}$$

   The Board makes their wage decision based upon these beliefs.

2. **Shareholder beliefs when they commit to signal threshold k_{st}**

   At the time that S commits to voting strategy, B has beliefs $\left(\mu_{bt|z_b}, \sigma_{bt|z_b}\right)$ and S has $\left(\mu_{at}, \sigma_a^2(\tau_t)\right)$ (the prior). S can discern $\left(\mu_{bt|z_b}, \sigma_{bt|z_b}\right)$ for any z_{bt}, which they integrate out int her objective function (18).

3. **Board and shareholder beliefs about shareholders' *ex ante* signal distribution at the time of the SOP vote**

   Before the SOP vote, when the shareholders commit to their threshold, both B and S know that the shareholders' aggregated signal will be

$$Z_{syt} = a + p\varepsilon_{st} + (1-p)\varepsilon_{yt} \tag{IA7}$$

with $Z_{syt} \sim N\left(\mu_{at}, \sigma_{at}^2 + \frac{\sigma_{z_s}^2 \sigma_y^2}{\sigma_{z_s}^2 + \sigma_y^2}\right)$.[2] Notice that shareholder beliefs about a do not update to $N\left(\mu_{bt|z_b}, \sigma_{bt|z_b}\right)$. This is because the timing convention in the model states that the wage w_t (and equivalently z_{bt}), productivity z_{yt} and the signal z_{st} are all revealed concurrently (see Figure IA1).

4. **Board and shareholder beliefs after the annual shareholder meeting**

   The 10-K and compensation committee report reveals the wage to shareholders, hence reveals z_{bt}. Shareholders vote at the annual meeting and Z_{syt} and thus z_{st} are revealed. Hence, B and S beliefs update to $\left(\mu_{at+1}, \sigma_a^2(\tau_t + 1)\right)$, by Prop. 2.

IA1.3. Assumptions about Shareholder strategy

See Section 3.3 for the full discussion of the Shareholder's strategy. I model this threat as an *ex ante* probability that the SOP will fail, which is increasing in the wage, and I make the following assumptions:

Assumption 1. *Shareholders commit to their voting strategy in advance of the annual shareholder meeting.*

I assume that S commits to an *ex ante* probability of vote failure to influence the Board's wage decision, even though *ex post* failure is costly. Because of the information revelation structure of the model (Figure IA1), S sets the probability of vote failure before seeing z_{bt}, z_{st}, or z_{yt}. The threat of vote failure does not need to be revealed to the Board before the annual shareholder meeting, however commitment forces S to play the threshold strategies that the Board expects. Unlike Kakhbod et al. (2023), there is no notion of cheap talk here. Commitment means S cannot choose an *ex ante* optimal non-zero failure probability and then renege at the shareholder meeting once the Board sets their wage.[3]

Assumption 2. *Shareholders are* myopic. *That is, the SOP vote is only influenced by today, and is not a fully dynamic problem.*

Effectively, this means that shareholders play a static game, while the Board plays a dynamic one. This assumption matches reality. There is ample evidence that voting in SOPs is influenced by short-run outcomes, such as current firm or stock performance (see Fisch et al., 2018; Novick, 2019, 2020).

[2]The variance of the signal is $Var(a + p\varepsilon_{st} + (1-p)\varepsilon_y) = \sigma_a^2 + p^2\sigma_{st}^2 + (1-p)^2\sigma_y^2$, where $p = \frac{\sigma_{z_s}^{-2}}{\sigma_{z_s}^{-2} + \sigma_y^{-2}}$. Expanding this expression out leads to the expression for the variance.

[3]Based on Assumption 1, Figure IA1 provides a more detailed version of the model timeline (slightly adapting Figure 2). In particular, in period 1 (or 1a and 1b), B and S set their strategies. These strategies are not revealed at this time, but this timing convention defines the notion of the Board's informational advantage. In particular, the Board plays their strategy *after* receiving signal; the shareholder plays their strategy *before*. the assumption of commitment forces S to stick with the strategy that B expects.

This makes the solution method much simpler, as I only need to solve for one value function in the numerical solution.

IA1.4. Full Derivation of Model Solution

Proposition 4. *The Board's problem can be written as*

$$V(\mu_a, \tau, w_-) = \max_{w(z_b,s)} \quad \exp\left(\mu_{b|z_b} + 0.5\left(\sigma^2_{b|z_b} + \sigma^2_y\right)\right) w(z_b, s)^\alpha - (1-\lambda_\tau) w(z_b, s) - \chi_B F^U_{\tilde{z}_s}(s \times w(z_b, s)) -$$
$$AC(w(z_b, s), w_-; \tau) + \delta_B \Big[f(\tau_t) V^R + (1 - f(\tau_t)) E_{b|z_b}\left[V(\mu'_a, \tau+1, w(z_b, s))\right]\Big] \quad \text{(IA8)}$$

where

- $\mu_{b|z_b}$ *and* $\sigma^2_{b|z_b}$ *are defined in* (IA5) *and* (IA6),
- $E_{b|z_b}[V] = F^U_{\tilde{z}_s}(s \times w)$ *where* $F^U_{\tilde{z}_s}$ *is the CDF as in* (16)
- $AC(w, w_-; \tau)$ *(adjustment cost) is defined in* (4),
- $f(\tau_t)$ *are CEO tenure-specific hazard rates, with* $f_0 = 0$ *and* $f(\tau_t) = 1$
- $V^R = V(\mu_0, 0, 0)$ *as in* (21)
- $\mu'_a \mid \mu_{b|z_b}, \tau \sim N\left(\mu_{b|z_b}, \sigma^2_{b|_{z_b}} - \sigma^2_a(\tau+1)\right)$ *from Prop.* 3

The shareholder's problem can be written as

$$\max_s \int_{z_b} f(z_b \mid \mu_a, \sigma^2_a) \left[\exp\left(\mu_a + 0.5\left(\sigma^2_{b|z_b} + \sigma^2_y\right)\right) w(z_b, s)^\alpha - w(z_b, s) - \chi_S CDF^U_s(s \times w(z_b, s)) \right] dz_b \quad \text{(IA9)}$$

where $f\left(z_b \mid \mu_a, \sigma^2_a\right)$ *is the density function of* z_b *given prior beliefs about CEO ability, and all other objects are defined as above.*

Proof. I start with (IA8). $\mu_{b|z_b}$ and $\sigma^2_{b|z_b}$ are Board beliefs after receiving their signal, hence are known from the perspective of the Board. As $A = \exp(a + \varepsilon_y)$, with a (and beliefs about a) normally distributed, I can write its expectation in terms of means and variances. The probability of vote failure is described in Section 3.3, but as brief overview it is given by the CDF $F^U_{\tilde{z}_s}$, of the unbiased (lognormal) wage of beliefs implied by realizations of $\tilde{z}_s$. The adjustment cost makes the Board's problem dynamic, as they have to factor in the effect of wages on the continuation value.

V^R is value if the CEO retires, so beliefs reset and there is no adjustment cost. In other words, the Board's problem reverts to its $t = 1$ value; it is constant for any state as the prior belief of ability about the CEO talent pool is distributed $N(\mu_0, \sigma_0^2)$ for any state.

The distribution of μ'_a conditional on μ_a and τ is given in Prop. 3. However, because the Board has beliefs $\left(\mu_{b|z_b}, \sigma^2_{b|z_b}\right)$, the variance of next period *mean* beliefs (not the variance of beliefs) at the time the Board makes their decision is $\sigma^2_{b|z_b} - \sigma^2_{a'}$. This quantity wil be used to take expectation over the continuation value. If the CEO continues, the tenure increases by 1, and the Board must consider the adjustment cost in the next period.

For (IA9), the objects are the same as the Board's problem, however the shareholders choose s after integrating out z_b. That is, shareholders figure out the Board's wage decision for each z_b, including how they would react to the choice of a particular s, and maximize expected operating income. Under Assumption 2, shareholders do not behave dynamically, and only vote on the current period.

The solution $(w(z_b), s)$ is to be found numerically, each $(w(z_b), s)$ is a best response under commitment (Assumption 1). To sketch the intuition of the solution, fix S' strategy s under commitment. The Board can then back out the probability of failure for each choice of $w(z_b)$, knowing that S must play the threshold. In other words, there is no notion of deviation for the Board. S just needs to maximize (IA9) for their strategy to be a best response; they cannot deviate at the vote and play a lower threshold. ■

IA1.5. Derivation of Model Statistics

This section derives several closed-form model statistics that are useful to interpret the magnitude of the main effects from the model. I can directly derive standard errors for closed-form functions of model parameters, which is useful for comparing across models.

SOP failure cost as a percentage of unbiased value. To interpret the magnitude of the SOP failure cost, I first develop a measure of unbiased firm value. Unbiased firm value is the discounted stream of future cash flows produced by the CEO if the CEO were paid the profit-maximizing wage, under the assumption that Board and shareholder beliefs remain fixed at (μ_0, σ_0^2). First, note that

$$w_0 = \arg\max_w E_0[A_0 w^\alpha - w] = \alpha^{\frac{1}{1-\alpha}} \times E_0[A_0]^{\frac{1}{1-\alpha}}$$

is the optimal unbiased wage, absent SOP. Using this, average (unbiased) firm value can be written as

$$\text{Average firm value} = V_0 = \sum_{t=1}^{\infty} \delta_B^t E_t[y_t - w_t]$$

$$= \sum_{t=1}^{\infty} \delta_B^t E_0[y_0 - w_0]$$

$$= \sum_{t=1}^{\infty} \delta_B^t E_0[A_0 \times (w_0)^\alpha - w_0]$$

$$= \frac{1}{1-\delta_B} \left[\exp\left(\mu_0 + 0.5\left(\sigma_0^2 + \sigma_y^2\right)\right) \times (w_0)^\alpha - w_0\right]$$

I can use unbiased firm value to interpret the magnitude of the SOP failure cost parameters χ_B and χ_S

$$\text{SOP failure cost (\% average value)} = \frac{\chi_{\{B,S\}}}{V_0} \tag{IA10}$$

Board capture as a share of surplus. To interpret the magnitude of my board capture parameter, I can express it in terms of how the Board decides to split up the surplus between the Board and shareholder. Focusing on the average CEO in the first year of tenure, and abstracting from dynamics and SOP, suppose the Board places the weight $v_\tau \in [0,1]$ on the CEO's utility (pure dollar wage), and $1 - v_\tau$ on company profits, so that their program is

$$\max_{w_\tau} \quad (1 - v_\tau) \times E_\tau([A] w_\tau^\alpha - w_\tau) + v_\tau \times w_\tau$$

As can be seen, this is the same as $\max_{w_\tau} E_\tau[A] w_\tau^\alpha - (1 - \lambda_\tau) w_\tau$, with $\lambda_\tau = \frac{v_\tau}{1-v_\tau}$, the main program from the paper. Equivalently defining $v_\tau = \frac{\lambda_\tau}{1+\lambda_\tau}$, there exists $\lambda_\tau^\star$ for which the CEO captures all profits. If $\lambda_\tau = 0$, the Board maximizes profits. The optimal split that the Board decides for the CEO is thus

$$\theta(\tau) = \frac{v_\tau w_\tau}{(1 - v_\tau) \times (E_\tau[A] w_\tau^\alpha - w_\tau) + v_\tau w_\tau} = \frac{\lambda_\tau w_\tau}{E_\tau[A] w_\tau^\alpha - (1 - \lambda_\tau) w_\tau}$$

The split $\theta(\tau)$ describes how much the Board tilts the surplus towards the CEO.

IA2. Estimation Appendix

IA2.1. Numerical Solution

The model requires 12 estimable parameters, the Board's discount rate and $\mathbf{T} + 1$ externally calibrated CEO separation rates. I externally calibrate the Board's discount factor $\delta_B = 0.9$, following Taylor (2010). The CEO separation rates are generated by calculating the cross-sectional proportion of CEOs that separate from their firm for a given tenure. I group the remaining 12 parameters as $\mathbb{P}$,

$$\mathbb{P} = \begin{bmatrix} \log \eta & \kappa & \sigma_0 & \sigma_y & \alpha & c_w & \sigma_{z_b} & \sigma_{z_s} & \lambda_0 & \lambda_1 & \chi_B & \chi_S \end{bmatrix} \tag{IA11}$$

The model's solution proceeds as such

1. Start with a given $\mathbb{P}$
2. Discretize each idiosyncratic shock into an N_z grid; e.g., fix the possible realizations of $\varepsilon_{z_{bt}}$, $\varepsilon_{z_{st}}$, $\varepsilon_{z_{yt}}$, while maintaining the same seed sequence across iterations. I set $N_z = 21$
3. Discretize the state space into a $\left(N_\mu, \mathbf{T}+1, N_w\right) = (21, 26, 15)$ grid, call it $\mathcal{S}$, where each tuple (μ_i, τ_j, w_k) indexes current mean belief about CEO ability, tenure (which fully determines beliefs of variance of CEO ability) and the previous wage.
4. Start with a guess of $V_0(\mu, \tau, w)$ as the solution to the static game (i.e., where there is no wage adjustment cost), so V_0 is just the Board's per-period expected utility given optimal choices. Each $V(\mu, \tau, w_{1:N_w})$ starts with the same value.
5. Use Gauss quadrature, interpolation and Prop. 2 to estimate the continuation value for each tuple (μ, τ, w)
6. Guess S' voting policy policy and compute $\mathcal{S}(\mu, \tau, w_-)$, or S' flow utility given the voting policy
7. Given S' policy, find B's wage policy $\omega(\mu, \tau, w_-)$ given S' policy using value function iteration on $V(\mu, \tau, w_-)$
8. Return to 5 and repeat until $\max|V_i - V_{i-1}| < \epsilon = 1\text{e}-5$ and $\max|\mathcal{S}_i - \mathcal{S}_{i-1}| < \epsilon = 1\text{e}-5$

This process returns the Board's wage policy for each element in in the state space and each realization on the grid of ε_{z_b}. Concurrently, it returns the shareholder's policy for each element in the state sapce.

IA2.2. Simulation

I set $N_f = 2000$ firms. Given $a \sim (0, \sigma_0)$, I draw a CEO of skill a for each firm. A CEO spell is the length of time the CEO is matched with a firm. Each period, for each firm, I generate realizations of $\varepsilon_{z_{bt}}$, $\varepsilon_{z_{st}}$ and ε_{yt} (which are drawn under the same sequence of quasi-random numbers each simulation iteration). Given the state, I use the policies described in Section IA2.1 to generate optimal choices. Beliefs update given realizations of $\varepsilon_{z_{bt}}$, $\varepsilon_{z_{st}}$ and ε_{yt}. At the end of each period, for CEOs with tenure $\tau > 0$, they separate (via firing, quitting or retirement) with exogenous probability f_τ.

I generate $N_S = 5$ samples for each simulation. I "fix" randomness across different simulations. That is, each $n_s \in N_S$ sample has the same seed across iterations, only the variance of each CEO ability and each shock changes. Within each simulation n_s, I allow 20 years of burn-on data, and use 15 years as my analysis sample.

IA2.3. Estimation

I estimate the 12 parameters in (IA11) As mentioned above, the Board's discount factor δ_B is calibrated to 0.9. (Taylor, 2010), and separation rates are calibrated to match observed separation rates in the sample. I estimate the remaining model parameters by finding a vector $\mathbb{P}$ of parameters that minimizes the weighted distance between simulated and data moments. That is, given model moment $m(\mathbb{P})$ and data moments $m(X)$ and an appropriate weighting matrix W, I minimize

$$\min_{\mathbb{P}} \left[m(\mathbb{P}) - m(X)\right]' W \left[m(\mathbb{P}) - m(X)\right] \tag{IA12}$$

I use the identity matrix as the weighting matrix. Given that some moments in my estimation are extremely precise, but secondary in terms of relevance for understanding the impact of SOP, using the optimal weighting matrix would prioritize precision over relevance. For example, the output residual is extremely precise, and while important for learning in the model, it does not have the same model-relevance as the observed SOP failure rate. I allow less precise but critical SOP-related moments to have equal impact on the estimation, which reduces efficiency at the gain of greater eocnomic relevance.

When calculating standard errors, I use the variance-covariance matrix of the data moments, calculated using the influence function approach (as such, parameter standard errors are larger than if estimated using the optimal weighting matrix). Table IA1 displays the moments, along with the model parameter each aims to identify (see Section 4.1).

IA2.4. Optimization algorithm

My goal is to find the global minimum of the SMM objective function described in Section IA2.3. To leverage the efficiency of parallel computing, I use a somewhat modified version of the TikTak global optimization algorithm described in Arnoud et al. (2019).[4] The modifications are designed to take advantage of high performance computing to minimize computing time. The global optimization routine can be described as such:

1. **Parallel local minimization**

[4] I modified code from https://github.com/tpapp/MultistartOptimization.jl, which is based upon the original TikTak code: https://github.com/serdarozkan/TikTak. See also Liu (2022) for a recent example.

Table IA1. Moment targeting exercise

This table displays the notation and description for each targeted moment, along with the parameter(s) it targets.

	Description	**Notation**	**Target**
(1)	SOP failure rate	s_0	χ_S
(2)	SOP fail wage sensitivity	s_1	$\{\chi_S, \chi_B\}$
(3)	SOP fail output sensitivity	s_2	$\{\sigma_{z_s}, \chi_S\}$
(4)	SOP failure residual variance	$Var(\epsilon^s)$	σ_{z_s}
(5)	Average log wage SOP pass	b_0	λ_0
(6)	Change in log wage SOP fail	b_1	$\{\chi_B, \chi_S\}$
(7)	Wage growth over tenure	b_2	λ_1
(8)	Wage residual variance	$Var(\epsilon^b)$	σ_{z_b}
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	c_w
(10)	Average log output	y_0	$\log \eta$
(11)	Elasticity of output to wage	y_1	α
(12)	Output residual variance	$Var(\epsilon^y)$	σ_y
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	σ_0
(14)	Average profit margin	Π / y	κ

i. Generate bounds for each parameter. This is a holistic step, yet the bounds should be narrow enough to allow for the subsequent quasi-random sequences to adequately cove the space, but wide enough so that I maximize the chance of finding the global minimum.

ii. Using the bounds, generate a Sobol sequence of length N. Sobol points are quasi-random points which are intended to mimic a draw from from a uniform distribution. In my setup, I set $N = 50,000$ or $N = 100,000$.

iii. For each $n \in N$ of the Sobol points, evaluate the objective function at n. Keep N_p of the points with the smallest initial function value. I set $N_p = 100$ or $N_p = 200$.

iv. For each $p \in N_p$, solve for the local minimum. In my setup, I use Nelder-Mead locally, which gives N_p “promising" candidates for the global minimum.

2. **Parallel global minimization**. This step slightly modifies the TikTak routine to take advantage of parallel computing. I employ SLURM with MPI to enable communication between ranked sets of iterations across the N_p points. This allows me to speed up the TikTak global optimization step, though at the expense of far greater expenditure of computing resources.

   i. Take the $p \in N_p$ candidates for the global minimum from above and sort in ascending order. Set $i = 1$, so the best minimum so far is indexed by i.

   ii. Take the best minimum so far, labeled p_i^*. Generate $N_p - i$ convex combinations using the TikTak methodology. That is, for $j \in N_p - i$, $p_{ji}^{cand} = \theta_{ji} p_i^* + (1 - \theta_{ji}) p_j$, where $\theta_{ji} \in [0, 1]$ and approaches 1 as j increases.

iii. Compute the local minimum of each $N_p - i$ point in parallel. If p_i^* is the best, then exit the routine and p_i^* is the candidate global minimum. Else,

iv. For the first j such that function value of p_{ji}^{cand} is less than that of p_i^*, stop all subsequent (unfinished) local minimization routines for $j' \in N_p - i$, and $j' > j$. Update i += p and return to ii.

This routine will return p_i^* as the global minimum.

3. **Polish global minimum**. Using stricter stopping criteria and a large number of function iterations, polish the global minimum p_i^* using a local minimization routine, i.e. Nelder-Mead.

IA2.5. Parameter Estimates and Model Fit for Subsample Analysis

Table IA2. Parameter estimates and model fit for subsample splits
This table displays the parameter estimates with standard errors and model fit for each subsample split presented in Table 7.

Panel A1. Low co-option

Moments

	Description	Notation	Data: Observed	Data: Simulated	Difference	*t*-stat
(1)	SOP failure rate	s_0	0.053	0.067	-0.014	-0.488
(2)	SOP fail wage sensitivity	s_1	0.057	0.060	-0.003	-0.077
(3)	SOP fail output sensitivity	s_2	-0.021	-0.039	0.019	0.411
(4)	SOP failure residual variance	$Var(\epsilon^s)$	0.032	0.045	-0.013	-1.905
(5)	Average log wage SOP pass	b_0	1.503	1.506	-0.003	-0.902
(6)	Change in log wage SOP fail	b_1	0.224	0.224	0.000	0.001
(7)	Wage growth over tenure	b_2	0.034	0.032	0.001	0.069
(8)	Wage residual variance	$Var(\epsilon^b)$	0.102	0.127	-0.025	-0.416
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	0.646	0.667	-0.021	-0.989
(10)	Average log output	y_0	7.651	7.674	-0.023	-0.227
(11)	Elasticity of output to wage	y_1	1.015	1.036	-0.022	-1.300
(12)	Output residual variance	$Var(\epsilon^y)$	2.105	2.121	-0.015	-0.871
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	1.964	1.957	0.007	0.838
(14)	Average profit margin	Π / y	0.173	0.218	-0.045	-1.807

Parameters

Description	Notation	Value
SOP failure costs		
Board SOP failure cost (% average firm value)	χ_B / V_0	4.28% (0.135%)
Shareholder SOP failure cost (% average firm value)	χ_S / V_0	1.49% (1.800%)
Estimated parameters		
Board SOP failure cost	χ_B	2.233 (0.0275)
Shareholder SOP failure cost	χ_S	0.776 (0.9640)
CEO board capture (constant)	λ_0	0.126 (0.0060)
CEO board capture (growth)	λ_1	0.082 (0.0052)
Prior std dev of CEO ability	σ_0	1.022 (0.0170)
Output—CEO wage elasticity	α	0.300 (0.0156)
Std dev of productivity shock	σ_y	1.358 (0.0126)
Scaling factor (output)	$\log \eta$	7.225 (0.1055)
Scaling factor (cost)	κ	0.340 (0.0096)
Std dev of Board signal	σ_{z_b}	0.646 (0.0028)
Std dev of Shareholder signal	σ_{z_s}	1.040 (0.0050)
CEO wage adjustment cost	c_w	5.395 (0.1147)

Panel A2. High co-option

Moments

			Data			
	Description	**Notation**	**Observed**	**Simulated**	**Difference**	***t*-stat**
(1)	SOP failure rate	s_0	0.080	0.035	0.044	1.840
(2)	SOP fail wage sensitivity	s_1	0.092	0.080	0.011	0.235
(3)	SOP fail output sensitivity	s_2	0.004	-0.004	0.008	0.095
(4)	SOP failure residual variance	$Var(\epsilon^s)$	0.049	0.016	0.033	5.433
(5)	Average log wage SOP pass	b_0	1.650	1.648	0.002	0.114
(6)	Change in log wage SOP fail	b_1	0.269	0.272	-0.003	-0.027
(7)	Wage growth over tenure	b_2	0.054	0.047	0.007	0.379
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.133	0.070	0.064	0.984
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	0.524	0.372	0.152	4.337
(10)	Average log output	y_0	7.586	7.775	-0.189	-1.223
(11)	Elasticity of output to wage	y_1	1.436	1.277	0.159	8.068
(12)	Output residual variance	$Var(\epsilon^y)$	3.059	2.954	0.104	4.409
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	2.744	2.767	-0.023	-1.565
(14)	Average profit margin	Π / y	0.243	0.279	-0.036	-0.750

Parameters

Description	**Notation**	**Value**
SOP failure costs		
Board SOP failure cost (% average firm value)	χ_B / V_0	1.33% (0.030%)
Shareholder SOP failure cost (% average firm value)	χ_S / V_0	1.58% (0.048%)
Estimated parameters		
Board SOP failure cost	χ_B	1.435 (0.0130)
Shareholder SOP failure cost	χ_S	1.712 (0.0033)
CEO board capture (constant)	λ_0	0.099 (0.0016)
CEO board capture (growth)	λ_1	0.105 (0.0027)
Prior std dev of CEO ability	σ_0	1.353 (0.0067)
Output—CEO wage elasticity	α	0.294 (0.0066)
Std dev of productivity shock	σ_y	1.460 (0.0102)
Scaling factor (output)	$\log \eta$	7.297 (0.1720)
Scaling factor (cost)	κ	0.351 (0.0030)
Std dev of Board signal	σ_{z_b}	0.441 (0.0090)
Std dev of Shareholder signal	σ_{z_s}	0.527 (0.0050)
CEO wage adjustment cost	c_w	6.002 (0.0538)

Panel B1. Low top five institutional ownership

Moments

	Description	Notation	Data Observed	Data Simulated	Difference	t-stat
(1)	SOP failure rate	s_0	0.062	0.064	-0.001	-0.069
(2)	SOP fail wage sensitivity	s_1	-0.017	-0.010	-0.008	-0.329
(3)	SOP fail output sensitivity	s_2	-0.076	-0.050	-0.026	-0.937
(4)	SOP failure residual variance	$Var(\epsilon^s)$	0.032	0.044	-0.012	-1.168
(5)	Average log wage SOP pass	b_0	1.767	1.762	0.005	0.484
(6)	Change in log wage SOP fail	b_1	0.220	0.109	0.110	2.107
(7)	Wage growth over tenure	b_2	0.046	0.006	0.040	4.864
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.100	0.138	-0.038	-0.833
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	0.674	0.623	0.050	2.297
(10)	Average log output	y_0	8.101	8.596	-0.494	-8.873
(11)	Elasticity of output to wage	y_1	1.221	1.135	0.086	14.833
(12)	Output residual variance	$Var(\epsilon^y)$	2.220	2.155	0.065	9.817
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	2.039	2.097	-0.058	-17.283
(14)	Average profit margin	Π / y	0.197	0.308	-0.110	-3.825

Parameters

Description	Notation	Value
SOP failure costs		
Board SOP failure cost (% average firm value)	χ_B / V_0	3.62% (0.295%)
Shareholder SOP failure cost (% average firm value)	χ_S / V_0	1.00% (0.014%)
Estimated parameters		
Board SOP failure cost	χ_B	2.327 (0.0638)
Shareholder SOP failure cost	χ_S	0.645 (0.0305)
CEO board capture (constant)	λ_0	0.069 (0.0181)
CEO board capture (growth)	λ_1	0.069 (0.0135)
Prior std dev of CEO ability	σ_0	0.997 (0.0240)
Output—CEO wage elasticity	α	0.365 (0.0414)
Std dev of productivity shock	σ_y	1.389 (0.0086)
Scaling factor (output)	$\log \eta$	7.958 (0.0927)
Scaling factor (cost)	κ	0.357 (0.0283)
Std dev of Board signal	σ_{z_b}	0.584 (0.0066)
Std dev of Shareholder signal	σ_{z_s}	1.451 (0.0051)
CEO wage adjustment cost	c_w	0.555 (0.0058)

Panel B2. High top five institutional ownership

Moments

	Description	Notation	Data Observed	Data Simulated	Difference	t-stat
(1)	SOP failure rate	s_0	0.072	0.072	0.001	0.044
(2)	SOP fail wage sensitivity	s_1	0.070	0.069	0.000	0.015
(3)	SOP fail output sensitivity	s_2	-0.006	-0.052	0.047	1.295
(4)	SOP failure residual variance	$Var(\epsilon^s)$	0.038	0.045	-0.007	-0.483
(5)	Average log wage SOP pass	b_0	1.444	1.446	-0.003	-0.946
(6)	Change in log wage SOP fail	b_1	0.230	0.226	0.004	0.073
(7)	Wage growth over tenure	b_2	0.052	0.054	-0.001	-0.127
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.112	0.101	0.011	0.323
(9)	Log wage autocovariance	$\gamma(w, w_{-})$	0.478	0.501	-0.023	-1.385
(10)	Average log output	y_0	7.332	7.396	-0.065	-1.033
(11)	Elasticity of output to wage	y_1	1.046	1.117	-0.071	-11.005
(12)	Output residual variance	$Var(\epsilon^y)$	2.560	2.514	0.046	8.063
(13)	CEO-average output variance	$Var(E_i[\tilde{y}])$	2.403	2.396	0.007	0.756
(14)	Average profit margin	Π / y	0.237	0.297	-0.060	-3.357

Parameters

Description	Notation	Value
SOP failure costs		
Board SOP failure cost (% average firm value)	χ_B / V_0	4.48% (0.057%)
Shareholder SOP failure cost (% average firm value)	χ_S / V_0	2.66% (0.062%)
Estimated parameters		
Board SOP failure cost	χ_B	3.107 (0.0109)
Shareholder SOP failure cost	χ_S	1.848 (0.0537)
CEO board capture (constant)	λ_0	0.185 (0.0028)
CEO board capture (growth)	λ_1	0.077 (0.0049)
Prior std dev of CEO ability	σ_0	1.238 (0.0050)
Output—CEO wage elasticity	α	0.278 (0.0026)
Std dev of productivity shock	σ_y	1.380 (0.0078)
Scaling factor (output)	$\log \eta$	6.997 (0.0671)
Scaling factor (cost)	κ	0.377 (0.0041)
Std dev of Board signal	σ_{z_b}	0.424 (0.0060)
Std dev of Shareholder signal	σ_{z_s}	1.918 (0.0126)
CEO wage adjustment cost	c_w	7.906 (0.0452)

References

Arnoud, A., Guvenen, F., and Kleineberg, T. (2019). Benchmarking global optimizers.

Fisch, J., Palia, D., and Solomon, S. D. (2018). Is say on pay all about pay: The impact of firm performance. *Harvard Business Law Review*, 8:101.

Kakhbod, A., Loginova, U., Malenko, A., and Malenko, N. (2023). Advising the management: A theory of shareholder engagement. *The Review of Financial Studies*, 36(4):1319–1363.

Liu, T. (2022). Bargaining with private equity: Implications for hospital prices and patient welfare.

Novick, B. (2019). Executive compensation: The role of public company shareholders. https://corpgov.law.harvard.edu/2019/07/31/executive-compensation-the-role-of-public-company-shareholders/ [Accessed: 2023-02-18].

Novick, B. (2020). "The goldilocks dilemma": A response to lucian bebchuk and scott hirst. In *Colum. L. Rev. Forum*, volume 120, page 80. HeinOnline.

Pastor, L. and Veronesi, P. (2009). Learning in financial markets. *Annual Review of Financial Economics*, 1(1):361–381.

Pinnington, J. (2023). Are passive investors biased voters?

Taylor, L. A. (2010). Why are CEOs rarely fired? Evidence from structural estimation. *The Journal of Finance*, 65(6):2051–2087.